

Extended object reconstruction in adaptive-optics imaging: the multiresolution approach

Roberto Baena Gallé^{1,2*}, Jorge Núñez^{1,2**}, and Szymon Gladysz³

¹ Departament d'Astronomia i Meteorologia i Institut de Ciències del Cosmos (ICC), Universitat de Barcelona (UB/IEEC), Martí i Franqués 1, E-08025 Barcelona, Spain.

² Observatori Fabra, Reial Acadèmia de Ciències i Arts de Barcelona, Camí de l'Observatori s/n, E-08002 Barcelona, Spain.

³ Fraunhofer-Institut für Optronik, Systemtechnik und Bildauswertung, Gutleuthausstraße 1, 76275 Ettlingen, Germany.
e-mail: szymon.gladysz@iosb.fraunhofer.de

Received 00 xxxxxx 2011 / Accepted 00 xxxxxx 2011

ABSTRACT

Aims. We propose the application of multiresolution transforms, such as wavelets (WT) and curvelets (CT), to the reconstruction of images of extended objects that have been acquired with adaptive optics (AO) systems. Such multichannel approaches normally make use of probabilistic tools in order to distinguish significant structures from noise and reconstruction residuals. Therefore, we also test the performance of two different probabilistic masks: one based on local correlation and the other on local standard deviation. Furthermore, we aim to check the historical assumption that image-reconstruction algorithms using static PSFs are not suitable for AO imaging.

Methods. We convolve an image of Saturn taken with the Hubble Space Telescope (HST) with AO PSFs from the 5-m Hale telescope at the Palomar Observatory and add both shot and readout noise. Subsequently, we apply different approaches to the blurred and noisy data in order to recover the original object. The approaches include multi-frame blind deconvolution (with the algorithm IDAC), myopic deconvolution with regularization (with MISTRAL) and wavelets- or curvelets-based static PSF deconvolution (AWMLE and ACMLE algorithms). We used the mean squared error (MSE) and the structural similarity index (SSIM), also known as the Wang-Bovik index, to compare the results. We discuss the strengths and weaknesses of the two metrics and the usefulness of using more than one metric to evaluate the results of imaging restoration.

Results. We found that CT produces better results than WT, as measured in terms of MSE and SSIM. In the wavelet domain, the mask based on local correlation provides better results than the mask based on standard deviation for the same number of iterations. Furthermore, multichannel deconvolution with a static PSF produces results which are generally better than the results obtained with the myopic/blind approaches (for the images we tested) thus showing that the ability of a method to suppress the noise and to track the underlying iterative process is just as critical as the capability of the myopic/blind approaches to update the PSF.

Key words. Instrumentation: adaptive optics - Techniques: image processing - Methods: miscellaneous

1. INTRODUCTION

The distortions introduced into images by the acquisition process in astronomical ground-based observations are well known. Apart from the most common, such as vignetting, non-zero background or bad pixels, which must be removed prior to any other analysis, atmospheric turbulence limits the spatial resolution of an image whereas the electronic devices used to acquire and amplify the signal introduces noise. An image is also corrupted by Poisson noise due to fluctuations in the number of received photons at each pixel (Andrews & Hunt 1977). The classical equation that describes the image formation process is:

$$\text{image} = [\text{PSF} * \text{object}] \diamond \text{noise} \quad (1)$$

where $*$ denotes the convolution operation. The symbol \diamond is a pixel-by-pixel operation which reduces to the simple addition in the case when noise is additive and independent of $[\text{PSF} * \text{object}]$, while for Poisson noise it is an operation which returns a random deviate drawn from a Poisson distribution with mean

equal to $[\text{PSF} * \text{object}]$. It is well known that direct inversion of equation 1 in the Fourier domain amplifies noisy frequencies close to the cut-off frequency. Hence, in the presence of noise, such a simple method cannot be used.

Several deconvolution approaches have been proposed in order to estimate the original signal from the seeing-limited and noise-degraded data. Since equation 1 is an ill-posed problem, with non-unique stable solutions, one approach is to regularize the Fourier inversion in order to constrain possible solutions (Tikhonov et al. 1987; Bertero & Boccacci 1998). This method generally imposes a trade-off between noise amplification and the desired resolution which generally leads to smooth solutions. Bayesian methodology (Molina et al. 2001; Starck et al. 2002b) allows a solution compatible with the statistical nature of the signal to be sought, leading to maximum likelihood estimators (MLE) (Richardson 1972; Lucy 1974) or maximum a posteriori (MAP) approaches if prior information is used, e.g., the positivity of the signal or entropy (Frieden 1978; Jaynes 1982).

All these methods can be enhanced through multichannel analysis by decomposing the signal in different planes, each of them representative of a certain scale of resolution. In such a decomposition, fine details in an image are confined to some planes, whereas coarse structures are confined to others. One

* e-mail: rbaena@am.ub.es

** e-mail: jorge@am.ub.es

of the most powerful ways to perform such decomposition is by means of the wavelet transform (WT). In particular in the astronomical context, the undecimated isotropic à trous algorithm (Holschneider et al. 1989; Shensa 1992; Starck & Murtagh 1994), also known as the Starlet transform, is often used. The WT creates a multiple representation of a signal, classifying its frequencies and, simultaneously, spatially localizing them in the field of view. For the specific case of the Starlet transform, this can be expressed by:

$$\mathbf{s} = \omega_0^s + \omega_1^s + \omega_2^s + \dots + \omega_N^s + \mathbf{r}_N^s \quad (2)$$

where \mathbf{s} is the signal being decomposed into wavelet coefficients, ω_j^s is the wavelet plane at resolution j and \mathbf{r}_N^s is the residual wavelet plane. We point out that equation 2 provides a prescription for direct reconstruction of the original image from all of the wavelet planes and the residual plane. The advantage of WT is that it allows for different strategies to be used for different wavelet planes, e.g., by defining thresholds of statistical significance.

Given the advantages of the multiresolution analysis the aforementioned deconvolution approaches have been adapted to work in the wavelet domain. The maximum-likelihood method (MLE) was modified into a two-channel algorithm (Lucy 1994) where the first channel corresponds to the signal contribution and the second to the background. The wavelet transform has also been applied to maximum-entropy deconvolution methods (Núñez & Llacer 1998) in order to segment the image and apply different regularization parameters to each region. Starck et al. (2001) generalizes the method of maximum-entropy within a wavelet framework, separating the problem into two stages: noise control in the image domain and smoothness in the object domain.

While WT has been widely used in astronomical image analysis and in deconvolution, the reported use (in the same context) of multi-transforms with properties that improve on or complement WT, is scarce. Such methods include, among many others: waveatoms (Demant & Ying 2007), which aim to represent signals by textures; and curvelets (Candès et al. 2006), which introduce orientation as a classification parameter together with frequency and position. Starck et al. (2002a) make use of the curvelet transform (CT) for Hubble Space Telescope (HST) image restoration from noisy data. In that paper, enhanced contrast was reported on the image of Saturn. Lambert et al. (2006) apply CT to choose significant coefficients from astroseismic observations while Starck et al. (2004) use CT for detection of non-Gaussian signatures in the observations of the cosmic microwave background. Since it is believed that CT is more suitable for representing elongated features such as lines or edges, one of the goals of this paper is to introduce CT into the Bayesian framework and show how it performs on images of planetary objects.

The aforementioned methods work with static PSFs, i.e., they do not update the PSF of an optical system which is supplied by the operator. Nevertheless, in ground-based imaging, whether with adaptive optics (AO) or without, there are always differences between science and calibration PSFs (Esslinger & Edmunds 1998). These differences may result from changes in seeing, wind speed, slowly varying aberrations due to gravity or thermal effects and, in AO imaging, from differential response of the wavefront sensor to fluxes received from the science and calibration objects. Quality of AO images is quantified with the Strehl ratio (SR) which is the ratio of the measured peak value of a point-source image to that of the diffraction-limited PSF, often given in percent. A perfect diffraction-limited image has SR of

100% while a seeing-degraded image on a large telescope can have SR lower than 1%. In this paper we will use SR to quantify the mismatch between the target and calibration PSFs. This commonly-occurring mismatch has prompted optical scientists, especially those working on AO systems, to investigate blind and myopic image restoration schemes (e.g. Lane 1992; Thiébaut & Conan 1995). A blind method works without any information about the PSF while a myopic approach relies on some initial PSF estimate which is then updated until a solution for both the object and the PSF is found.

Conan et al. (1999) compared their myopic approach to a basic, “unsupervised” Richardson-Lucy scheme and found, not surprisingly, that the myopic deconvolution is more stable. Nevertheless, to our knowledge there have not been any thorough and fully-fledged efforts to compare the performance of modern, static-PSF approaches to blind/myopic methods in the context of AO imaging, and so the preference for the latter algorithms still has to be justified. Our goal is to partially fill this gap. At this stage we want to mention that this article is a companion paper to Baena Gallé & Gladysz (2011) where we showed that a modern, static-PSF code is capable of extracting accurate differential photometry from AO images of binary stars even when given a mismatched PSF. In the current paper we extend the analysis to a more complex object and we analyze the effect of noise as well as that of the mismatched PSF. We show the importance of noise control, specifically how advanced noise suppression can set off the lack of PSF-update capability in the case of very noisy observations. More generally, we want to hint at the opportunities for exchange of ideas between the communities preferring myopic and static-PSF approaches.

The paper is organized as follows: in Section 2, the algorithms AWMLE, ACMLE, MISTRAL and IDAC are described. These algorithms represent different philosophies with regards to the deconvolution problem. AWMLE and ACMLE perform a classical static-PSF deconvolution within the wavelet or curvelet domain, MISTRAL is intended for myopic use, and IDAC can be used as a blind or a myopic algorithm. Section 3 describes the dataset we used and how we applied each of the four algorithms. A description of the mean squared error (MSE) and the structural similarity index (SSIM) used to compare the reconstructed images can be found in section 4. Section 5 presents the comparison of two strategies of multiresolution support. The performances of the wavelet and the curvelet transforms are also discussed. Section 5 also contains the performance comparison of the four algorithms mentioned above. Section 6 summarizes and concludes the paper.

2. DESCRIPTION OF THE ALGORITHMS

This section is not intended to describe the algorithms in detail but rather to offer a brief overview of their characteristics and their historical uses and performances. We would also like to point out that in the text to follow we do not keep the nomenclature originally used by the respective authors. We make the following attempt at standardization: the object, or unknown, is represented with \mathbf{o} , the image or dataset is \mathbf{i} , and the PSF is \mathbf{h} . Upper-case notation is used for Fourier representation. The two-dimensional pixel index is r while f is spatial frequency. Index of a wavelet/curvelet scale is represented with j while a single frame in a multi-frame approach is indexed with i . Finally, the symbol $\hat{\cdot}$ is used for estimates.

2.1. AWMLE

The Adaptive Wavelet Maximum Likelihood Estimator (AWMLE) is a Richardson-Lucy-type algorithm (Richardson 1972; Lucy 1974) that maximizes the likelihood between the dataset and the projection of a possible solution onto the data domain, considering a combination of the Poissonian shot noise, intrinsic to the signal, and the Gaussian readout noise of the detector. This maximization is performed in the wavelet domain. The decomposition of the signal into several channels allows for various strategies to be used depending on a particular channel's scale. This is a direct consequence of the fact that in a WT decomposition the noise, together with the finest structures of the signal will be transferred into the high-frequency channels while coarse structures will be transferred into the low-frequency channels.

The general mathematical expression that describes AWMLE is:

$$\hat{\mathbf{o}}^{(n+1)} = K \hat{\mathbf{o}}^{(n)} \left[\mathbf{h}^- * \frac{\sum_j \left(\omega_j^{\mathbf{o}_h^{(n)}} + \mathbf{m}_j (\omega_j^{\mathbf{i}'} - \omega_j^{\mathbf{o}_h^{(n)}}) \right)}{\mathbf{h}^+ * \hat{\mathbf{o}}^{(n)}} \right] \quad (3)$$

where \mathbf{o} is the object to be estimated, \mathbf{h}^+ is the point spread function (PSF) which projects the information from the object domain to the image domain, while \mathbf{h}^- is its reverse version performing the inverse operation, i.e., the conversion from the image domain to the object domain. The so-called direct projection \mathbf{o}_h is an image of the object projected onto the data domain by means of the PSF at iteration n , it appears in the numerator of equation 3 already decomposed in its wavelet representation.

The parameter $K = \frac{\mathbf{i}}{\hat{\mathbf{o}}^{(n+1)}}$ is a constant to conserve the energy.

The variable \mathbf{i}' is a modified version of the dataset which appears here because of the explicit inclusion of the readout Gaussian noise into the Richardson-Lucy scheme (Núñez & Llacer 1993). This represents a pixel-by-pixel filtering operation in which the original dataset is substituted by this modified version. It must be calculated for each new iteration by means of the following expression:

$$i'(r) = \frac{\sum_{k=0}^{\infty} (ke^{-(k-i(r))^2/2\sigma^2} [(o_h(r))^k / k!])}{\sum_{k=0}^{\infty} (e^{-(k-i(r))^2/2\sigma^2} [(o_h(r))^k / k!])} \quad (4)$$

Where \mathbf{i} is the original dataset, r is the pixel index, σ is the standard deviation of the readout noise and \mathbf{o}_h is the direct projection of the object onto the data domain. In the absence of noise ($\sigma \rightarrow 0$) the exponentials in equation 4 are dominant when $k = i(r)$ and then $i'(r) \rightarrow i(r)$, i.e., the modified dataset converges to the original one. We note that by its definition, \mathbf{i}' is always positive while \mathbf{i} can be negative due to the presence of readout noise.

In expression 3, the symbols $\omega_j^{\mathbf{o}_h^{(n)}}$ and $\omega_j^{\mathbf{i}'}$ correspond to wavelet coefficients, in channel j , of the multiscale representation of the direct projection $\mathbf{o}_h^{(n)}$ as well as that of the modified dataset \mathbf{i}' , respectively.

Of particular importance is the probabilistic mask \mathbf{m}_j , centered on pixel r , which locally determines whether a significant structure is present or not in channel j . If the probability of finding a source within the mask is considered to be high then \mathbf{m}_j will be close to 1, and so the dataset \mathbf{i}' will be the only remaining term in the numerator and will be compared, at iteration n , with the estimated object at that iteration $\mathbf{o}^{(n)}$. If the presence of a

signal within the mask is considered to be insignificant, then \mathbf{m}_j will decrease to zero and the object will be compared with itself (by means of its projection $\mathbf{o}_h^{(n)}$), so the main fraction in equation 3 will be set to unity (at channel j in the vicinity of pixel r) and this would stop the iterative process. Therefore, the probabilistic mask \mathbf{m}_j is used to effectively halt the iterations in a Richardson-Lucy scheme. In other words, it allows a user to choose from a wide range of the maximum number of iterations to be executed, where the reconstructions corresponding to these different numbers will not exhibit significant differences between them, which arise in the classic Richardson-Lucy algorithm due to the amplification of noise.

AWMLE in the wavelet domain was first introduced by Otazu (2001). Baena Gallé & Gladysz (2011) presented the algorithm in the context of AO imaging and obtained differential photometry in simulated AO observations of binary systems. The accuracy of aperture photometry performed on the deconvolution residuals was compared with the accuracy of PSF-fitting, a classic approach to the problem of overlapping PSFs from point sources (Diolaiti et al. 2000). It was proven that AWMLE yields similar, and often better, photometric precision compared to StarFinder independently of the stars' separation. Even though AWMLE does not update the PSF as it performs deconvolution, it was shown that the resulting photometric precision is robust to mismatches between the science and the calibration PSF up to 6% in terms of difference in SR which was the maximum tested mismatch.

2.2. ACMLE

Equation 2 offers a direct reconstruction formula for the wavelet domain, i.e., the sum of all the wavelet planes (and the residual one) in which the original image had been decomposed allows one to obtain back that same original image. This explains the presence of summation in the numerator of expression 3, where a combination of different wavelet coefficients, some belonging to the direct projection of the object (\mathbf{o}_h) and some belonging to the modified dataset (\mathbf{i}'), creates the correction term in this Richardson-Lucy scheme.

In the curvelet domain, the reconstruction formula does not have this simple expression. For both: the forward and the inverse transforms, double Fourier inversions and complex operations with arrays are required (Candès et al. 2006). Nevertheless, there is an intuitive way to obtain a kind of *à trous curvelet transform*, a method to recover the original image by using a simple sum as it is done in equation 2. This would consist of decomposing the image into N curvelet scales, setting to zero all the coefficients except those corresponding to a certain scale, and applying the inverse curvelet transform to transfer the remaining coefficients to the spatial domain. One by one, the same operation would be targeted at the other of curvelet scales.

This would create N images in the spatial domain, each of them representative of a certain curvelet scale. All of them could be added in order to recover the original image and, consequently, expression 3 would still be valid. Unfortunately, this relatively straightforward and intuitive way to introduce CT into AWMLE equation invalidates most of the advantages of directly working in the curvelet domain. Firstly, it requires one forward transform plus N inverse transforms at each iteration. Secondly, the sparsity concept, which allows one to work with only a few non-zero coefficients and is the basis of many of the multiresolution transforms, is abandoned in this naive approach. As mentioned by Baena Gallé & Núñez (2010), CT tends to separate the

smallest signal structures from the noise in the spatial domain better than WT. However, it is impossible to separate the signal from noise completely, and so the SNR of the finest curvelet plane is lower than the SNR of the finest wavelet plane, since less signal is represented in such high-frequency plane in the spatial domain (SNR here refers to the ratio of intensity between the remnants of objects in a particular plane and the finest structures created by noise during the classification process). In other words, the identification of the correct information is faster and more successful in the curvelet domain, especially in the highest frequency scales, when compared to the use of the representation of each curvelet scale in the spatial domain.

Given the above reservations about the usage of CT in the spatial domain we decided to work with the following expression:

$$\hat{\mathbf{o}}^{(n+1)} = K\hat{\mathbf{o}}^{(n)} \left[\mathbf{h}^- * \frac{\mathbb{C}^{-1}(\mathbb{C}(\mathbf{o}_h^{(n)}) + \mathbf{m}_j(\mathbb{C}(\mathbf{i}') - \mathbb{C}(\mathbf{o}_h^{(n)})))}{\mathbf{h}^+ * \hat{\mathbf{o}}^{(n)}} \right] \quad (5)$$

where \mathbb{C} and \mathbb{C}^{-1} mean, respectively, the forward and the inverse curvelet transforms. The mask \mathbf{m}_j is now calculated in the curvelet domain and combines coefficients from \mathbf{o}_h and \mathbf{i}' in order to create a new curvelet correction term which is inversely transformed to the spatial domain and compared with the object estimate at each iteration. The curvelet transform used in equation 5 corresponds to the so-called second generation CT and is implemented in the software CurveLab¹. This particular implementation of the CT exhibits a robust structure based on a mother curvelet function of only three parameters (scale, location and orientation), which is faster and simpler to use in comparison with the first generation CT based on a complex seven-index structure, which relies on the combined usage of the Starlet and the ridgelet transforms (Starck et al. 2010).

2.3. MISTRAL

The Myopic Iterative STep-preserving Restoration ALgorithm (MISTRAL) is a deconvolution method within the Bayesian framework that jointly estimates the PSF and the object using some prior information about both these unknowns (Mugnier et al. 2004). This joint Maximum A Posteriori (MAP) estimator is based on the following expression:

$$[\hat{\mathbf{o}}, \hat{\mathbf{h}}] = \arg \max [p(\mathbf{i}|\mathbf{o}, \mathbf{h}) \times p(\mathbf{o}) \times p(\mathbf{h})] = \arg \min [J_i(\mathbf{o}, \mathbf{h}) + J_o(\mathbf{o}) + J_h(\mathbf{h})] \quad (6)$$

where $J_i(\mathbf{o}, \mathbf{h}) = -\ln p(\mathbf{i}|\mathbf{o}, \mathbf{h})$ is the joint negative log-likelihood that expresses fidelity of the model to the data (\mathbf{i}), $J_o(\mathbf{o}) = -\ln p(\mathbf{o})$ is the regularization term, which introduces some prior knowledge about the object (\mathbf{o}) and $J_h(\mathbf{h}) = -\ln p(\mathbf{h})$ accounts for some partial knowledge about the PSF (\mathbf{h}). The symbol p in the above expressions corresponds to the probability density function of a particular variable.

MISTRAL does not use separate models for the Poisson and the readout components of noise. Instead, a nonstationary Gaussian model for the noise is adopted. What this means is that a least-squares optimization with locally-varying noise variance is employed:

$$J_i(\mathbf{o}, \mathbf{h}) = \sum_r \frac{1}{2\sigma^2(r)} [i(r) - (o * h)(r)]^2 \quad (7)$$

where r stands for pixel index. This prior makes it easier to compute the solution with gradient-based techniques as compared to the Poissonian likelihood which contains a logarithm. The Gaussian assumption is typical (Andrews & Hunt 1977) and it can be considered a very good approximation for bright regions of the image. The assumption can cause problems for low-light-level data recorded with modern CCDs of almost negligible readout noise.

The prior probability, $J_o(\mathbf{o})$, is modelled to account for objects which are a mix of sharp edges and smooth areas such as those that we deal with in this article. The adopted expression contains an edge-preserving prior that is quadratic for small gradients and linear for large ones. The quadratic part ensures a good smoothing of the small gradients (i.e., of noise), and the linear behaviour cancels the penalization of large gradients (i.e., of edges). Such combined priors are commonly called $L_2 - L_1$ (Green 1990; Bouman & Sauer 1993). The $L_2 - L_1$ prior adopted in MISTRAL has the following expression:

$$J_o(\mathbf{o}) = \mu\delta^2 \sum_r \phi(\nabla o(r)/\delta) \quad (8)$$

where $\phi(x) = |x| - \ln(1 + |x|)$ and where $\nabla o(r) = [\nabla_x o^2(r) + \nabla_y o^2(r)]^{1/2}$. Here, $\nabla_x \mathbf{o}$ and $\nabla_y \mathbf{o}$ are the object finite-difference gradients along x and y , respectively. Equation 8 is effectively $L_2 - L_1$ since it adopts the form $\phi(x) \approx x^2/2$ when x is close to 0 and $\phi(x)/|x| \rightarrow 1$ when x tends to infinity. The global factor μ and the threshold δ are two hyperparameters that must be adjusted by hand according to the level of noise and the object structure. Some strides towards semi-automatic setting of these parameters were made by Blanco & Mugnier (2011).

The regularization term for the PSF, which introduces the myopic criterion into equation 6, assumes that the PSF is a multidimensional Gaussian random variable. This assumption is justified when one deals with long exposures which are, by definition, sums of large numbers of short-exposures. As such they are Gaussian by the Central Limit Theorem. Adopting these conditions, $J_h(\mathbf{h})$ has the form:

$$J_h(\mathbf{h}) = \frac{1}{2} \sum_f \frac{|H(f) - H_m(f)|^2}{E[|H(f) - H_m(f)|]^2} \quad (9)$$

This prior is expressed in the Fourier domain whereby uppercase notation denotes Fourier transformation. The term $H_m(f) = E[H]$ is the mean transfer function and $E[|H(f) - H_m(f)|]^2$ is the associated power spectral density (PSD) with f denoting spatial frequency. This Fourier-domain prior bears some resemblance to equation 7. Indeed, it also assumes Gaussian statistics and it draws the solution, in the least-squares fashion, towards the user-supplied mean PSF while obeying the error bars given by the PSD (which give the variance at each frequency). The PSF prior leads to band-limitedness of the PSF estimate because the ensemble average in the denominator, $E[|H(f) - H_m(f)|]^2$, should be zero above the diffraction-imposed cut-off.

In practice equation 9 relies on the availability of several PSF measurements. The mean PSF and its PSD are estimated by replacing the expected values ($E[.]$) by an average computed on a PSF sample. When such a sample of several PSFs is not available, as we have assumed in our work, then \mathbf{H}_m is made equal to the Fourier transform of the single supplied PSF, and $E[|H(f) - H_m(f)|]^2$ is computed as the circular mean of $|\mathbf{H}_m|^2$. These relatively large error bars are intentional: they account for the lack of knowledge about the PSD when given only a single PSF measurement.

¹ <http://www.curvelet.org>

In the original MISTRAL paper (Mugnier et al. 2004) the code was presented mainly in the context of planetary images, for which the object prior (equation 8) was developed. The experimental data presented therein was obtained on several AO systems and covered a wide range of celestial objects such as Jupiters satellites Io and Ganymede and the planets Neptune and Uranus. MISTRAL was applied to the study of the asteroids Vesta (Zellner et al. 2005) and 216-Kleopatra (Hestroffer et al. 2002), and was also used to monitor surface variations on Pluto over a 20-year period (Storrs & Eney 2010).

2.4. IDAC

Multi-Frame Blind Deconvolution (MFBD) (Schulz 1993; Jefferies & Christou 1993) is an image reconstruction method relying on the availability of several images of an object. In addition, many of the MFBD algorithms rely on short exposures. This was originally dictated by the notion that in imaging through turbulence short-exposure images contain diffraction-limited information while the long exposures do not (Labeyrie 1970). Before the advent of AO the only way to obtain diffraction-limited data from the ground was to record short exposures which could then be processed by one of several “speckle imaging” methods (Knox & Thompson 1974; Lohmann et al. 1983). Therefore, MFBD was originally proposed in the context of speckle imaging.

The MFBD code we use in this paper is called IDAC, for Iterative Deconvolution Algorithm in C^2 , and it is an extension of the Iterative Blind Deconvolution (IBD) algorithm proposed by Lane (1992). IDAC performs deconvolution by numerically minimizing a functional which is composed of four constraints:

$$\epsilon = E_{im} + E_{conv} + E_{bl} + E_{Fm} \quad (10)$$

where

$$E_{im} = \sum_{r \in \gamma} [\hat{o}(r)]^2 + \sum_{i=1}^M \sum_{r \in \gamma} [\hat{h}_i(r)]^2 \quad (11)$$

is the image domain error which penalizes the presence of negative pixels (γ) in both the object (\mathbf{o}) and the PSF (\mathbf{h}) estimates. The subscript i refers to an individual data frame.

The so-called convolution error is:

$$E_{conv} = \frac{1}{N^2} \sum_{i=1}^M \sum_f |I_i(f) - \hat{O}(f)\hat{h}_i(f)|^2 B_i(f) \quad (12)$$

which quantifies the fidelity of the reconstruction (\hat{O}) to the data (\mathbf{I}) in the Fourier domain. The term B_i is a binary mask that penalizes frequencies beyond the diffraction-imposed cut-off, N^2 is a normalization constant where N is the number of pixels in the image.

The third constraint is called the PSF band-limit error and it is defined as:

$$E_{bl} = \frac{1}{N^2} \sum_{i=1}^M \sum_f |\hat{h}_i(f)|^2 B'_i(f) \quad (13)$$

It prevents the PSF estimates from converging to a δ function and the object estimate from converging to the observed data. The term B'_i is a binary mask that is unity for spatial frequencies greater than 1.39 times the cut-off frequency and zero elsewhere.

The last constraint is the Fourier modulus error:

$$E_{Fm} = \frac{1}{N^2} \sum_f [|\hat{O}(f)| - |O_e(f)|]^2 \Phi(f) \quad (14)$$

where O_e is a first estimate of the Fourier modulus of the object obtained, e.g., from Labeyrie’s speckle interferometry method (Labeyrie 1970) and Φ is a signal-to-noise filter. Jefferies & Christou (1993) show with several simulations that this constraint is especially important since it incorporates relatively high SNR information from a complete dataset formed by many frames. On the other hand, Fourier modulus of an object is only recoverable directly if one has a set of speckle images of an unresolved source to be used as a reference. There are workarounds to this problem, most notably reference-less approaches due to Worden et al. (1977) and von der Lühe (1984), but these solutions are applicable only to non-compensated imaging and could not be used in our tests.

While MFBD algorithms were initially proposed in the context of speckle imaging, there is nothing preventing their application to long-exposure AO images which now contain diffraction-limited frequencies. In principle there are some inherent advantages of working with M frames instead of using a single, co-added long exposure. The multi-frame approach reduces the ratio of unknowns to measurements from 2 : 1 in single-image blind deconvolution to $M + 1 : M$ in multi-frame deconvolution. On top of the PSF band-limit constraint (equation 13) concurrent processing of many frames means that the PSF cannot converge to the δ function (this would have yielded an object equal to the data, but the data is generally temporally variable while the object is assumed constant in MFBD, which is a good assumption in the context of astronomical imaging on short time-scales). MFBD algorithms are very successful in the case of strongly varying PSFs so that the target is easily distinguished from the PSFs. On the other hand, the goal of AO is to stabilize the PSF. This implies less PSF diversity from one observation to another so that other constraints become more useful.

The code IDAC can be regarded as a precursor but also a representative of a wider class of MFBD algorithms. We will mention here the PCID code (Matson et al. 2009) which has the capability to estimate the PSFs either pixel by pixel in the image domain or in terms of a Zernike-based expansion of the phase in the pupil of the telescope. It has been shown that such a PSF re-parameterization leads to object estimates which are less noisy and have higher spatial resolution (Matson & Haji 2007).

IDAC was applied to speckle observations of the binary system Gliese 914 and in the process the secondary component was resolved into two stars (Jefferies & Christou 1993). Additionally, IDAC was one of the five codes used in our study of photometric accuracy of image reconstruction algorithms (Gladysz et al. 2010).

3. DATASET DESCRIPTION AND METHODOLOGY

3.1. Dataset description

The image we used to test the algorithms corresponds to an observation performed with the fourth detector of the WFPC2 camera (Trauger et al. 1994) (the so-called planetary camera -PC-, with a pixel size of 0.046'') installed on the HST. This is a picture of Saturn (the whole planet is visible) with a dynamic range of up to 975 counts. In order to obtain well-defined edges and transitions between the object and the background, the image of Saturn was preprocessed so that all the pixels below a certain

² <http://cfao.ucolick.org/software/idac/>

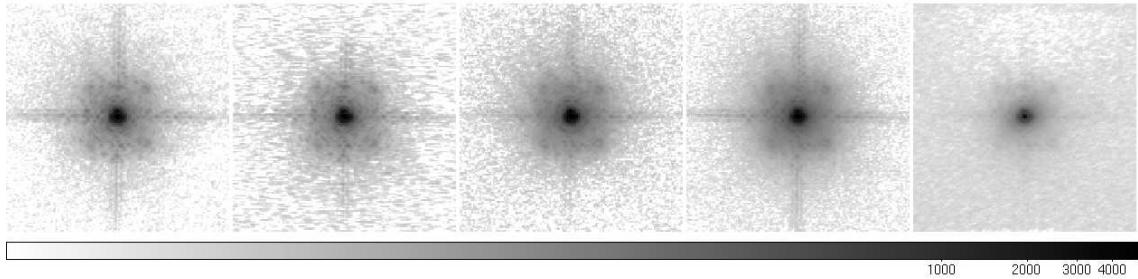


Fig. 1. Far left: science PSF with the Strehl ratio (SR) = 53%. Middle left: science PSF with SR = 51%. Center: reference PSF with SR = 53%. Middle right: reference PSF with SR = 45%. Far right: reference PSF with SR = 36%. Left and middle left PSFs were used to blur the HST image of Saturn. The other three PSFs were used as reference PSFs for the algorithms. Logarithmic scale, inverted colors.

Table 1. PSFs used for the simulated observations

Star	SR	exp. time (ms)	
PSF1	53%	1416	Science PSF
—	51%	—	Science PSF
PSF2	53%	2832	Reference PSF
—	45%	—	Reference PSF
PSF3	36%	1416	Reference PSF

threshold were set to zero, thus enhancing the visibility of the Cassini and Encke divisions. This image was considered a representation of the true object.

The “ground truth” image was subsequently convolved with an AO PSF. For tests of AWMLE and ACMLE with the local-correlation mask (see Section 3.2) it was mandatory to have two blurred images (we chose PSFs with Strehl ratios equal to 51 and 53%, see Table 1). This is because in these approaches two images are deconvolved simultaneously (Fig. 2.B).

While the σ -based approach (see Section 3.2) does not in principle require more than one image the same two images were used again for consistency. For the same reason MISTRAL was applied to two images, instead of one, and the two separate reconstructions were averaged. In the case of MFBD restorations with IDAC we used only the 51%-Strehl-ratio PSFs to blur the image(s). We want to remark at this stage that this requirement to have more than one image for AWMLE and ACMLE with the local-correlation mask does not fundamentally limit their applicability for astronomical imagery. A double image of the same object could be achieved in a real situation in several ways: two different observations on two consecutive nights, two outputs of a camera equipped with a beam splitter, by dividing a number of short-exposure frames into two different datasets or via the thinning method (Llacer et al. 1993), which allows an image to be split into two halves each of which preserves the Poissonian and Gaussian statistical nature of the original.

The PSFs were obtained with the 1024x1024 PHARO infrared camera (Hayward et al. 2001) on the 5-m Hale telescope at the Palomar Observatory (Troy et al. 2000). Closed-loop images of single stars were recorded using the $0.040'' \text{ pixel}^{-1}$ mode. The images were cropped to size 150×150 pixels which corresponds to a field of view (FOV) of $6''$. The observations were acquired in the K band ($2.2\mu\text{m}$) where the diffraction limit is $0.086''$ so that the data meet Nyquist-sampling requirements. The filter of the observations was Brackett Gamma (BrG) and each of the PSFs in Table 1 corresponds to a sum of 200 frames, each with an exposure time being either 1416 or 2832 ms. The individual frames were registered via iterative Fourier shifting to produce shift-and-add images (Baena Gallé & Gladysz 2011).

The angular size on the sky over which the AO PSF can be assumed to be almost spatially invariant is the so-called isoplanatic angle. This parameter becomes larger at longer wavelengths. As specified by Hayward et al. (2001) the isoplanatic angle is approximately $50''$ in K-band at Palomar. The angular size of the Saturn image is $20.5''$ so we can assume the PSF remains constant throughout the FOV. The difference in pixel scales between the Saturn image and the PSFs from Palomar is very small ($0.046''$ vs. $0.040''$) and therefore we did not re-bin the images to match their pixel scales. The paper is devoted to comparison of image restoration algorithm and not to the performance evaluation of the AO system at Palomar.

We used IDAC in the multi-frame mode. One of the goals of our work was to check the trade-off between the diversity provided by more frames vs. SNR per frame. One can think of this as an exposure time optimization. For a constant total observation time per object one can use shorter exposures and hope to exploit PSF variability in subsequent image restoration with MFBD, or opt instead to use a smaller number of images with better SNR per frame. Therefore, out of the original 200-frame dataset we have produced datasets of 10, 20, 50 and 100 binned PSFs. These binned PSFs, together with the original 200-frame dataset, were used as blurring kernels for the Saturn image. For AWMLE, ACMLE and MISTRAL, which are single-image restoration codes, we only used the summed PSF. All PSFs were normalized to have total power equal to unity before the convolution procedure.

The blurred observations were subsequently corrupted with noise. We explored three levels of noise: only readout noise of standard deviation $\sigma = 1$ or $\sigma = 10$, and shot noise plus readout noise of level $\sigma = 10$. These levels correspond to noise which was added to the summed images (two images for AWMLE, ACMLE and MISTRAL). For MFBD with IDAC we worked with several images and decided to add to each frame the amount of noise which would result in the same SNR per summed image had the images been summed and *then* processed as in the case of the other three algorithms. This means adding noise with levels: $1 \cdot \sqrt{50}$, $10 \cdot \sqrt{50}$ and Poisson noise plus $10 \cdot \sqrt{50}$ to the 50-frame dataset, for example. This way we wanted to test the benefits of exploiting PSF diversity vs. higher noise per frame, as mentioned earlier.

One of the goals of the project was to test the susceptibility of the algorithms to mismatch between the science and calibration PSFs. Therefore, we used a matched PSF (SR=53%), a mismatched PSF (SR=45%) and a highly-mismatched PSF (SR=36%) as inputs to AWMLE, ACMLE, MISTRAL and MFBD. The last SR value corresponds to half of the maximum SR in K band (63%) predicted in simulations for the AO system in Palomar (Hayward et al. 2001). Table 1 lists the PSFs used

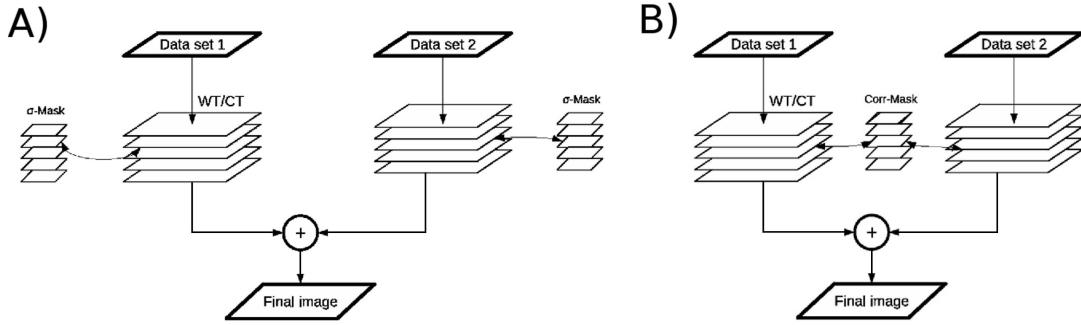


Fig. 2. A: A flow graph illustrating the execution of the AWMLE algorithm with the σ -based mask. Datasets are decomposed by WT and each of the resulting planes is *independently* analyzed with the σ -based mask. B: A respective graph for execution with the correlation-based mask. After decomposition the planes are *simultaneously* analyzed with the correlation-based mask.

in the tests. Differences in SR for the same star arise because of the changing seeing. All five PSFs used in the simulations are presented in Figure 1.

3.2. Methodology

The four algorithms are very different and require different strategies to obtain the final result. For AWMLE and ACMLE, it is convenient (although not mandatory) to perform an initial decomposition of the dataset, either in the wavelet domain or in the curvelet domain. This can give the user an idea as to the number of planes to use during the deconvolution process. Additionally, such a preliminary decomposition could inform the user whether it is sensible to perform deconvolution in the highest-frequency plane where the finest structures together with noise have been classified. When working in such planes there is a trade-off between the benefit of recovering information from the finest wavelet/curvelet plane and the undesirable reconstruction of non-significant structures. Both transforms, WT and CT can also be combined into a dictionary of coefficients (Fadili & Starck 2006), although for the sake of simplicity we execute them independently.

The main difficulty when using a Richardson-Lucy-type algorithm is determining at which iteration to stop the deconvolution process. In AWMLE and ACMLE, this problem is solved by probabilistic masks. The mask can apply significance thresholds to a given location in a given plane to selectively deconvolve statistically similar regions. This concept is called multiresolution support (Starck et al. 2002b). Probabilistic masks are used to stop the deconvolution process automatically in parts of the image where significant structure cannot be discerned. We use two different kinds of probabilistic mask. The first is based on the local standard deviation within the mask itself:

$$m_\sigma = \begin{cases} 1 - \exp \left\{ \frac{-\frac{3}{2}(\sigma_i - \sigma_\chi)^2}{2\sigma_\chi^2} \right\} & \text{if } \sigma_i - \sigma_\chi > 0 \\ 0 & \text{if } \sigma_i - \sigma_\chi \leq 0 \end{cases} \quad (15)$$

with:

$$\sigma_i = \sqrt{\frac{\sum_{p \in \Phi} (\chi_{j,p})^2}{n_f}}$$

where σ_i is the standard deviation within the window Φ centered on pixel i of the plane χ_j (pixels within that window are indexed with p), n_f is the number of pixels contained in the window and σ_χ is the global standard deviation in the corresponding wavelet or curvelet plane. This last value can be approximated by decomposing an artificial Gaussian noise image, with standard deviation equal to that in the dataset, into wavelet or curvelet coefficients.

The second mask is based on the local correlation between two images which are deconvolved simultaneously (but independently). If we have two observations of the same object, it can be assumed that the information pertaining to that object remains constant, while structures due to noise or atmospheric turbulence introduce some scatter in the information. The correlation mask measures, at each iteration, the similarity between the same region in two different images and, if noise amplification or speckle reconstruction is detected (in the sense that the local correlation is reduced), the algorithm is stopped for that region. The expression for that mask is:

$$m_c = \begin{cases} \left(\frac{\sigma_{ij}}{\sigma_i \sigma_j} \right)^\alpha & \text{if } 0 < \sigma_{ij} \leq 1 \\ 0 & \text{if } -1 < \sigma_{ij} \leq 0 \end{cases} \quad (16)$$

where σ_i and σ_j are again standard deviations within windows Φ centered on the pixel i of the plane χ_i of each image (both windows have the same position so $i = j$) and σ_{ij} is the local covariance within such a window. Parameter α penalizes the fact that the correlation tends to 1 very quickly. We found that a relatively high value, $\alpha = 24.5$, was enough to slow this down. One can see that the same correlation mask is applied to both images, which therefore undergo the same number of effective iterations across the FOV.

There is no need to work exclusively with one of the two masks and reject the other. The standard-deviation mask and the correlation mask can be combined in the deconvolution process. In general a good strategy is to use the standard-deviation mask within the highest-frequency wavelet/curvelet planes, where small significant structures or point-like sources can stand out above the noise, and to use the correlation mask

for the other planes, where distortions in large signal structures introduced by speckles can be partially corrected.

Figure 2 shows the flow diagram for AWMLE when either of the two masks is used. Standard-deviation masks are applied independently to either of the images (Fig. 2.A) while in the correlation-based approach the same mask is applied to both images simultaneously (Fig. 2.B). Using the correlation mask does not in itself imply a combined deconvolution of both images. The two threads in Figure 2.B are executed independently but, at each iteration, they make use of the same values to classify significant coefficients locally in each plane. The threads in Figure 2.A are also executed independently but the probabilistic masks used for either of the images can take different values. In both cases the two outputs of the deconvolution process are then averaged in order to obtain the final reconstruction.

CurveLab software, which is used to introduce the CT into ACMLE, can perform digital curvelet decomposition via two different implementations. The first one is the so-called USFFT digital CT and the second one is known as the digital CT via wrapping. Both differ in the way they handle the grid which is used to calculate the FFT in order to obtain the curvelet coefficients. This grid is not defined in typical Cartesian coordinates but rather it emulates a polar representation, more suited to the mathematical framework defined by Candès et al. (2006). The USFFT version has the drawback of being computationally more intensive with respect to the wrapping version, since the latter makes a simpler choice of the grid to compute the curvelets (Starck et al. 2010). Hence, for the reason of computational efficiency, the wrapping version was used in ACMLE.

The user has to provide CurveLab with the values of some parameters. Apart from the number of curvelet scales, the number of orientations or angles of representation in the second scale is also required as input from the user. This parameter will automatically set the number of angles for the rest of the scales. Evidently, the higher the number of orientations the longer the algorithm will need to perform CT, the higher the overall redundancy and the higher the computational cost for calculating the probabilistic masks at each scale. We decided to set this parameter to 16 as a trade-off between having a complete representation of all the possible orientations present in the image and the total execution time. Finally, a third parameter decides if the curvelet transform is replaced by an orthogonal wavelet representation at the final scale, i.e., the highest-frequency scale. This translates to less redundancy and a faster execution of the algorithm. Since we did not see significant differences in preliminary tests when choosing either one type of representation or the other, we set this parameter to zero, i.e., the final scale was represented by means of wavelets.

MISTRAL performs the minimization process of equation 6 by the partial conjugate-gradient method. The code requires that the user provides values for the hyperparameters μ and δ (see equation 8) which balance smoothing imposed by too much regularization against noise amplification. This has to be done by trial and error although the authors of the algorithm provide some suggestions in the original MISTRAL paper, namely that μ be set to around unity and δ be set to the norm of the image gradient ($\|\nabla i\| = [\sum_p |\nabla i(p)|^2]^{1/2}$). In practice the user experiments with various values of these parameters, centered on the values suggested by Mugnier et al. (2004), and chooses the image reconstruction which is most visually appealing. In our tests we found that $\mu = 10$ and $\delta = 2$ yielded the best results. We let the code run for the maximum number of iterations set to 10^6 . It usually converged after 3000 iterations.

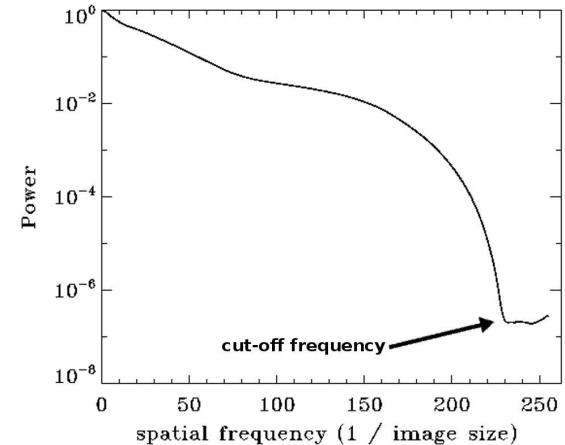


Fig. 3. Normalized power spectrum of the 53% SR PSF (Fig. 1 middle left). The diffraction- and noise-limited cut-off frequency was determined to be approximately 230 pixels. The PSF was embedded in an array of zeros to match the size of the science image (512×512 pixels).

IDAC requires the user to provide the value of the diffraction-limit cut-off. Caution should be taken here: the code requires that all the supplied images be of the same dimensions. Therefore, when one has a PSF image which is smaller in size than the target image, then the PSF should be embedded in an array of zeros. Subsequently, the diffraction-limit cut-off should be estimated from the zero-padded, and not the original, Fourier-transformed PSF (see Figure 3). Another parameter that should in theory affect the restorations, a scalar quantifying users confidence in the supplied PSF, was found to have negligible effect on the outputs.

As with MISTRAL we let the code run for the maximum number of iterations set to 10^6 . For images with high level of noise ($\sigma = 10$ and Poisson noise plus $\sigma = 10$) it converged quickly, after 15-20 iterations, which can be considered a good behaviour as noise did not become amplified. For cases with low noise ($\sigma = 1$) the algorithm converged after 50-100 iterations and produced generally sharper reconstructions.

4. Image quality metrics

One of the goals of this work was to compare the quality of reconstructions yielded by several codes. Therefore, metrics were needed for image quality assessment. Here we have chosen the mean squared error (MSE) and the so-called structural similarity index (SSIM) (Wang et al. 2004). The former is probably the most commonly used metric in image processing (e.g. Mugnier et al. 2004). On the other hand, SSIM was developed to mimic some of the hypothesized properties of the human visual system (HVS) and has been shown to overcome some of the weaknesses of the MSE.

4.1. Structural Similarity Index

The principal idea underlying the structural similarity approach is that the HVS is adapted to extract *structural* information from images, and therefore, a measurement of structural similarity (or distortion) should provide a good approximation to perceptual image quality (as obtained from human observers). The metric, SSIM, is a function of two images, denoted here by \mathbf{x} and \mathbf{y} . If \mathbf{x} is the reference image, then SSIM can be regarded as a mea-

surement of the quality of the second image \mathbf{y} . SSIM is defined as follows:

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \left[\frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \right]^\gamma \quad (17)$$

where μ_x and μ_y are the mean values of the images \mathbf{x} and \mathbf{y} , respectively. Parameters σ_x and σ_y are the standard deviations and σ_{xy} is the covariance between the two images. The metric is actually a multiplication of three functions that compare, respectively, the luminance, contrast and structure of the two images; i.e., $\text{SSIM}(x, y) = [l(x, y)^{\alpha_1} \cdot c(x, y)^{\alpha_2} \cdot s(x, y)^{\alpha_3}]$. These functions have the following expressions:

$$l(x, y) = \frac{(2\mu_x\mu_y + C_1)}{(\mu_x^2 + \mu_y^2 + C_1)} \quad (18)$$

$$c(x, y) = \frac{(2\sigma_x\sigma_y + C_2)}{(\sigma_x^2 + \sigma_y^2 + C_2)} \quad (19)$$

$$s(x, y) = \frac{(\sigma_{xy} + C_2/2)}{(\sigma_x + \sigma_y + C_2/2)} \quad (20)$$

The functions can be weighted in order to adjust the relative importance of each of them. For the sake of simplicity, we always have $\alpha_1 = \alpha_2 = \alpha_3 = 1$, which yields expression 17. It is worth to mention that these three functions are relatively independent, which is physically logical since a change in luminance and/or contrast of the imaged scene should not impact the perception of the structure of an object.

The SSIM has the properties of: symmetry, i.e., $\text{SSIM}(x, y) = \text{SSIM}(y, x)$; boundedness, whereby $\text{SSIM}(x, y) = 1$ indicates that the two images are identical; and uniqueness, in the sense that the previous upper bound is only reached when $x = y$. Note that $\text{SSIM}(x, y)$ can take negative values if x and y have an inverse correlation between them, i.e., $\sigma_{xy} < C_2/2$.

Parameters C_1 and C_2 are intended to ensure the results are stable when either $(\mu_x^2 + \mu_y^2)$ or $(\sigma_x^2 + \sigma_y^2)$ are close to zero or when $(\mu_x \cdot \mu_y \cdot \sigma_{xy})$ is identically zero, as is the case for pixels belonging to the background in the original image. The inclusion of background pixels between Saturn's body and the innermost ring, or in the vicinity of the outermost ring, is motivated by the behaviours of the deconvolution algorithms in these regions, as will be discussed later. Additionally, including background pixels gives a measurement of the ability of an algorithm to remove from the background (and reintroduce in the object) the light scattered by the PSF.

The choice of C_1 and C_2 is somehow arbitrary but, as commented by Wang et al. (2004), the performance of the SSIM index is almost insensitive to variations of these parameters. We decided to set C_1 to 300 and C_2 to 30 in order to assign more weight to those pixels that belong to the planet, rather than to the background. Nevertheless, in the initial tests we have not encountered substantial differences in the appearance of SSIM maps for different values of both C_1 and C_2 .

As suggested by its proponents, it is preferable to apply SSIM locally, by means of sliding windows centered on a given pixel, rather than over the entire image, in order to visualize space-variant features or distortions (Wang & Bovik 2006). Furthermore, such quality maps can provide more information about the local degradation of the image. Parameter γ in equation 17 is used to enhance the visibility of such maps. In any

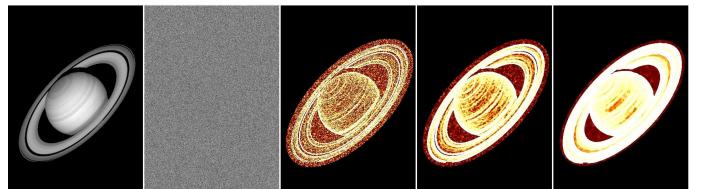


Fig. 4. Far left: original image. Middle left: white Gaussian noise ($\sigma = 10$) or absolute error map of the distorted image. Center: SSIM map of the distorted image created with 2×2 sliding windows. Middle right: SSIM map of the distorted image created with 4×4 sliding windows. Far right: SSIM map of the distorted image created with 8×8 sliding windows. Higher brightness indicates better quality. For a given position the structural similarity index was calculated only if all the pixels in a window belonged to the object and/or to the background in the vicinity of the object.

case, it is also possible to compute a single overall quality measure of the entire image by averaging all values calculated at the individual local windows. This average is called the mean structural similarity index (MSSIM):

$$\text{MSSIM}(x, y) = \frac{1}{M} \sum_{j=1}^M \text{SSIM}(x_j, y_j) \quad (21)$$

Where M is the number of local windows. At this point, we want to remark that the MSSIM metric must not be viewed as an “absolute truth” about the quality of the different reconstructions. Since Saturn image has zones with very different characteristics (e.g., background, main body, rings and planetary poles) MSSIM should be seen as a rough overview of an algorithm’s performance, but an inspection of the SSIM maps and reconstructions themselves is essential as well.

Figure 4 shows the SSIM maps produced for the Saturn image after adding white Gaussian noise. The reader may note that SSIM indicates worse quality for the smooth parts of the object, for example the main body of Saturn, than for the edges. This attribute of SSIM correlates well with perceptual image quality which is known to depend on the so-called contrast masking effect (Wang & Bovik 2006). What this means is that the visibility of the noise-corrupted signal, as perceived by human observers, varies significantly at different locations for objects which are a combination of smooth areas, textures and sharp edges. This can be quantified using a space-variant metric. On the other hand, a global metric such as MSE, even when applied locally, would yield the same results for various parts of the image shown in Figure 4. Of course, this ability of SSIM to capture local image quality depends on the size of the sliding window, as can be seen by comparing the three rightmost panels of Figure 4.

We note that SSIM is a widely-used metric for testing image quality in the remote sensing applications (González Audícan et al. 2005; Otazu et al. 2005).

4.2. Mean Squared Error

The mean squared error is probably the oldest and most often used metric to evaluate the resemblance of two images. It is defined as:

$$\text{MSE} = \frac{1}{M} \sum_{j=1}^M (x_j - y_j)^2 \quad (22)$$

Where, in this context, M is the total number of pixels of the object and in the surrounding area. It is also very common

to use the related metric, the peak signal-to-noise ratio (PSNR), defined by:

$$\text{PSNR} = 10 \log_{10} \frac{L^2}{MSE} \quad (23)$$

Where L is the dynamic range of the image. These metrics are very easy to compute and measure the pixel-by-pixel departure between the reconstruction and the reference object. The MSE and PSNR have a clear physical meaning: they quantify the energy of the error. Hence, they are well suited to the task of estimating absolute photometric error between the two images.

However, Wang & Bovik (2006) show that MSE exhibits poor correlation with the real (as perceived by humans) image quality, e.g., important modifications of a reference image, like defocusing or Gaussian noise contamination, produce relatively stable MSE values, whereas small spatial shifts or rotations, while having negligible effect on the subjective image quality, can yield very different MSEs or PSNRs. In the context of the Richardson-Lucy approach, where noise is amplified after a certain number of iterations, MSE and PSNR metrics can be insensitive to such amplification. Large visual distortions produced by noise, that might invalidate the scientific benefit of the reconstruction, could appear transparent for MSE because pixel-by-pixel absolute departure from the true value has not to be necessarily very high.

On the other hand, SSIM measures the local resemblance between the deconvolved image and the reference image by means of variance and correlation terms. Therefore, it is very sensitive to the presence of noise and distortions produced by defocusing, which are the typical distortions present in the deconvolution problem we are dealing with in this paper. On the other hand, we have detected the weakness of SSIM in detecting luminance changes (photometric differences). Although in principle it should provide a measurement of the photometric accuracy of reconstructions by means of equation 18, in practice, relatively large departures between both images do not necessarily produce poor SSIM values if such images show the same visual structure or trend in their respective numerical ranges.

Summarizing, MSE and SSIM metrics can be considered complementary to each other. Therefore, usage of both metrics constitutes a sound methodology for evaluating the performance of different deconvolution algorithms. The former quantifies photometric/radiometric departures more accurately, while the latter is more sensitive to the loss of perceivable image quality which can be the result of, e.g., noise amplification or residual blur.

In general, we have preferred to use PSNR rather than MSE to quantify differences between images, since the former has the same tendency as SSIM, i.e., higher values correspond to better reconstructions. This common property facilitates easy side-by-side comparisons. We set $L = 975$ in equation 23 to match the dynamic range of the “ground truth” object. PSNR maps were generated using expressions 22 and 23 but on a pixel-by-pixel basis in order to show the behaviour of this metric locally.

5. RESULTS AND DISCUSSION

5.1. Correlation-based mask versus σ -based mask

The tests of the two types of probabilistic masks were conducted using AWMLE in the wavelet domain, which was the framework used by Otazu (2001) to test his original design, i.e., AWMLE with σ -based mask. The algorithm was applied to two images, corresponding to the 51 and 53% SR, and corrupted with white

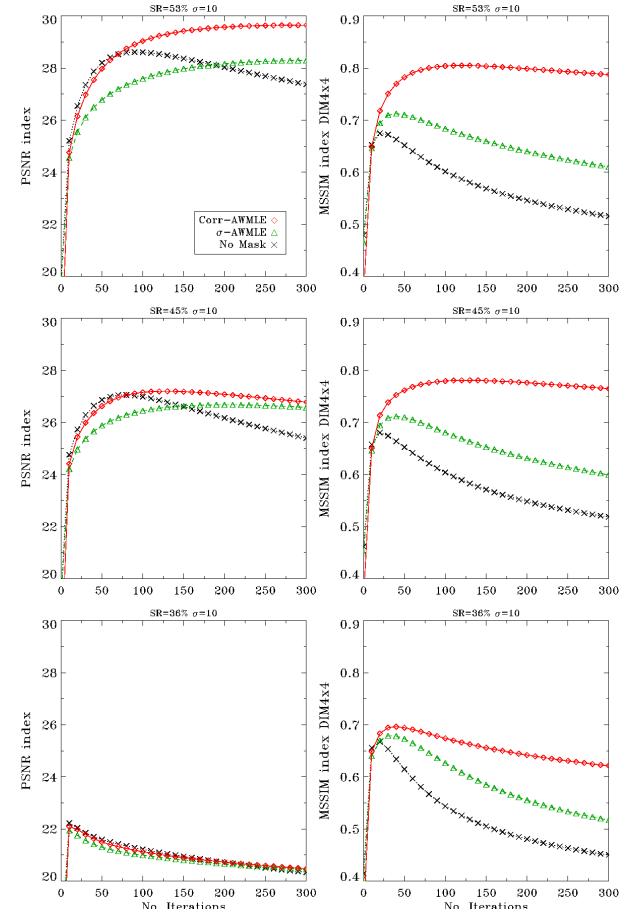


Fig. 5. Top row: evolution of the PSNR and MSSIM with respect to the number of iterations for the Saturn image degraded and corrupted with Gaussian noise ($\sigma = 10$) and using the SR=53% reference PSF. The left column shows PSNR index. The right column shows MSSIM index measured within 4x4 sliding windows. Middle row: SR=45% reference PSF. Bottom row: SR=36% reference PSF. Diamonds: AWMLE with correlation-based mask. Triangles: AWMLE with standard deviation-based mask. Crosses: AWMLE with no mask.

Gaussian noise of $\sigma = 10$. Both images were decomposed into three wavelet planes plus a wavelet residual. As mentioned before, three levels of mismatch between the science and calibration PSF were analysed. The wavelet planes were also deconvolved without using any kind of mask which, in practice, results in the execution of the Richardson-Lucy approach albeit in the wavelet domain.

The algorithm was executed for 300 iterations, storing intermediate results every 10 iterations. All the outputs were analyzed by means of PSNR and SSIM, using 4x4 sliding windows. SSIM was computed only in those windows where all the pixels belonged to the object or a small border around it (for the appreciation of how big this border is see Fig. 4). The results from such windows were averaged in order to obtain the global evolution of SSIM with respect to the number of iterations (Fig. 5).

As can be seen in figure 5, both masks were able to control degradation in PSNR, typical to the MLE approach, across a wide range of the number of iterations. When no mask is used the general quality of reconstruction quickly drops. Although Otazu (2001) assumed that the σ -based mask is able to stop noise amplification, one can observe (Fig. 5, right column) that this resistance is not permanent and eventually image quality

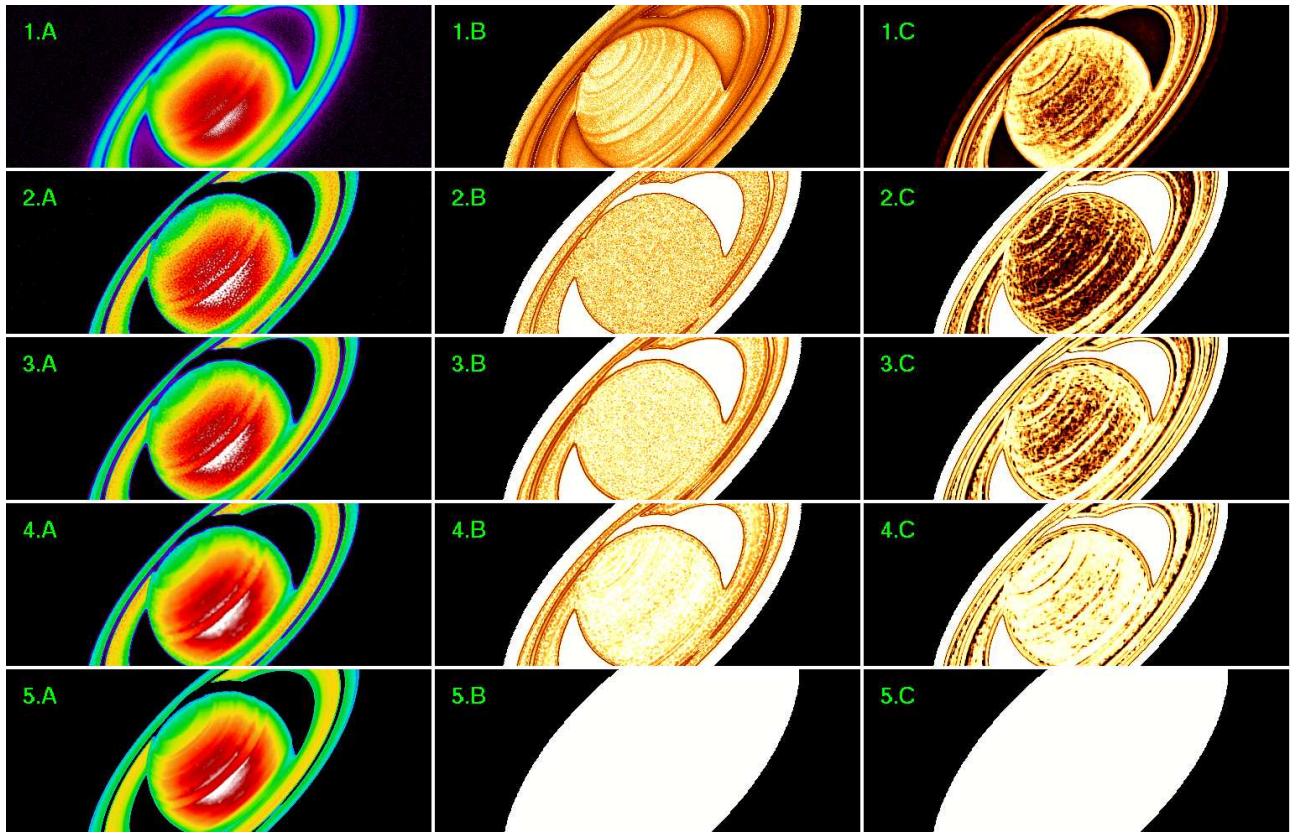


Fig. 6. 1.A: Degraded and corrupted image with Gaussian noise ($\sigma = 10$). 2.A: Reconstruction in the wavelet domain with no mask. 3.A: Reconstruction in the wavelet space with mask based on the local standard deviation. 4.A: Reconstruction in the wavelet domain with mask based on the local correlation. 5.A: Original image. Column B: PSNR maps. Column C: 4x4 SSIM maps. All reconstructions used the SR=45% PSF as the calibrator. The reconstructions were stopped at 100 iterations. Column A is represented on a linear scale from 0 to 975, matching the dynamic range of the original image, i.e., subpanel 5.A. Column B is represented on a linear scale from 0 to 50. Column C is represented on a linear scale from 0 to 1. In columns B and C higher brightness indicates better quality.

can drop significantly when this mask is employed in deconvolution. It is interesting to note that photometric accuracy of the reconstruction with the σ -based mask remains almost constant while image quality gets degraded by noise (Fig. 5, row 2 and Fig. 6, subpanel 3.A), thus showing the different behaviour and performance of PSNR and SSIM metrics. The correlation mask was able to stabilize the influence of noise more effectively. Furthermore, PSNR and MSSIM were higher for the correlation mask than for the σ -based mask. For example, at iteration number 100 and when the matched 53% SR PSF was used, the difference in MSSIM was $\sim 15\%$ and the “gap” between both masks is a constant of 1.5dB in PSNR. Comparing the evolutions of PSNR and MSSIM (Fig. 5, first row, no mask employed) one can see that when the reconstruction achieves its maximum photometric resemblance to the “ground truth”, at iteration ~ 100 , amplification of the noise is already so evident that the reconstruction shows a big degradation which eventually invalidates it (Fig. 6, row 2).

It should also be mentioned that a mismatch of $\sim 7\%$ in SR between the science PSF and the reference PSF (SR=45%) does not lead to large differences in the results (Fig. 5, middle row), with respect to those obtained with a well-matched reference PSF (SR=53%, Fig. 5, top row). The usage of the SR=36% reference PSF, which implies a mismatching of $\sim 16\%$ SR, can be seen to have a more noticeable effect on the process of noise control (Fig. 5, bottom row).

Figure 6 shows PSNR and 4x4 SSIM maps of the reconstructed images, as well as the reconstructions themselves. Results of the Richardson-Lucy method (in the sense that no mask was used, Fig. 6, subpanel 2.A) produced even worse SSIM maps than those obtained from the original, blurred and noise-corrupted image. This is due to visible noise amplification despite the resolution enhancement achieved from the reconstruction. The use of the masks clearly improves the results, more so if the correlation-based mask is applied in the wavelet domain (Fig. 6, row 4).

5.2. Wavelet vs. curvelet transform

The dataset described in Subsection 5.1 was also used to compare the quality of reconstructions when either the wavelet (WT) or the curvelet (CT) representation of the object was used. The schemes presented in equations 3 and 5 were executed in both domains. Again, the same two types of masks were employed: one based on the standard deviation and another based on the local correlation. Hundred iterations was considered to be enough to enhance the resulting image without the noise reconstruction affecting its general quality.

Looking at figure 7, at first glance there is no great difference between the wavelet and curvelet reconstructions. The CT and the WT produce similar PSNR and 4x4 SSIM maps within the body of the planet (Fig. 7, subpanels 2.B/C and 3.B/C) whereas WT has reconstructed the inner ring slightly better (subpanels

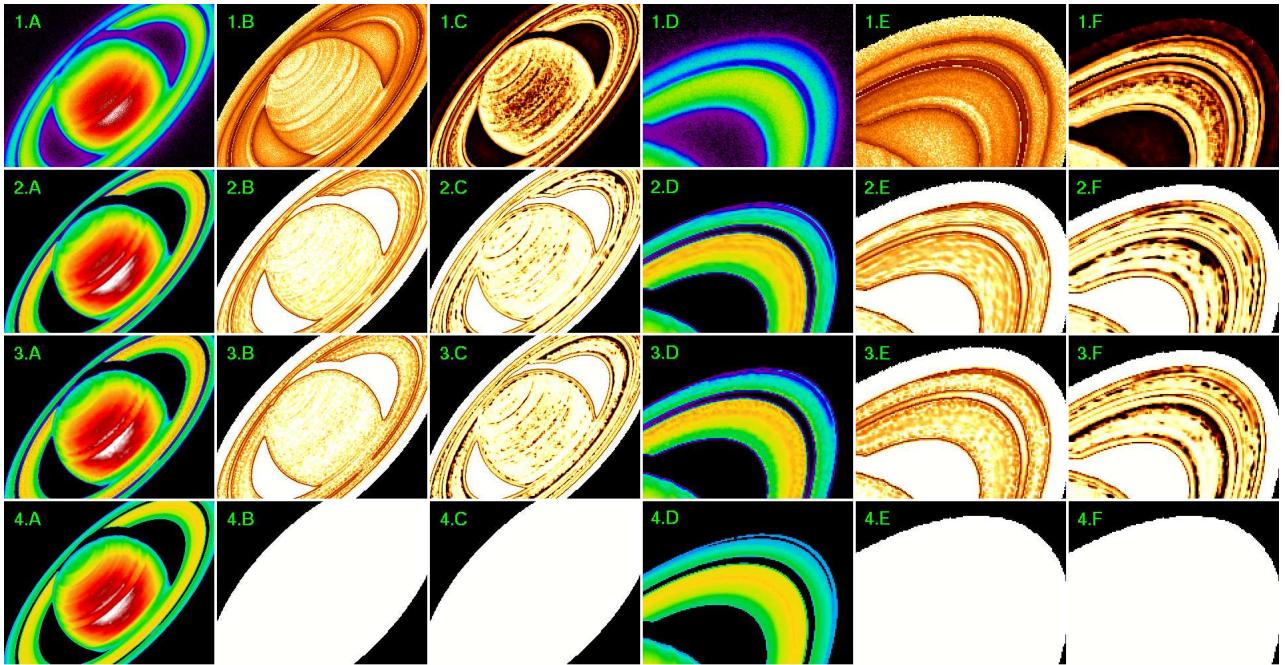


Fig. 7. Row 1: Degraded and corrupted image with Gaussian noise ($\sigma = 10$). Row 2: Reconstruction in curvelet domain with mask based on the local standard deviation. Row 3: Reconstruction in wavelet domain with mask based on the local correlation. Row 4: Original image. Columns B and E: PSNR maps. Columns C and F: 4×4 SSIM maps. Column D: fragment of the image in column A. All reconstructions used the SR=45% PSF as the calibrator. The reconstructions were stopped at 100 iterations. Columns A and D are represented on a linear scale from 0 to 975, matching the dynamic range of the original image, i.e., subpanel 4.A. Columns B and E are represented on a linear scale from 0 to 50. Columns C and F are represented on a linear scale from 0 to 1. In columns B, C, E and F higher brightness indicates better quality.

2.E/F and 3.E/F). On the other hand, CT has brought out some elongated features present in the original image on the inner ring which are now more pronounced than in the original (subpanels 2.D and 4.D). Looking directly at the reconstructions (Fig. 7, subpanels 2.A and 3.A) one can see that CT has reconstructed lines and visual transitions more homogeneously.

In order to determine which of the two transforms performs better as a whole, the mean value and standard deviation averaged across all the PSNR and SSIM maps (4×4 windows) were computed for two noise levels ($\sigma = 1$ and $\sigma = 10$ readout noise). The results are presented in Figure 8.

Apart from the logical conclusions related to the general improvement corresponding to low-level noise and well-matched PSF, it is possible to observe that CT works better than WT on this dataset in most of the situations. If σ -based mask is used, CT always exhibits an improvement with respect to WT. When PSF with SR=45% is used as calibrator, this improvement can be measured to be around 1dB in terms of PSNR or 5-12% in terms of MSSIM depending on the noise level, thus showing that resistance to noise amplification is more robust in the curvelet domain.

As in section 5.1, correlation and σ -based masks were also applied in the curvelet domain. On this occasion, we did not find large differences between the results. However, some artifacts or elongated structures were visible in the correlation mask results, especially in the presence of shot plus readout noise. Such artifacts are related with the result of wrong identifications of curvelet coefficients or non-homogenous reconstruction of different angle orientations at the same curvelet scale. This suggests that the correlation mask, in the curvelet domain, cannot be generalized for all the scales, as we do in the wavelet domain, but should be modified for each scale and each orientation.

The σ -based based mask in the curvelet domain produces slightly better results than the correlation-based mask in the wavelet domain. The slight decrease in reconstruction quality that can be observed between the standard-deviation and the correlation-based masks when CT is used, is related to the appearance of artifacts in the reconstructions, as was already mentioned.

5.3. Static-PSF approach versus blind/myopic approaches

One of the goals of the current work was to test the static-PSF reconstruction approach (based on multiresolution support) vs. typical methods used in AO (based on the myopic/blind philosophy). Again, we used three levels of PSF mismatch as well as three noise levels: Gaussian $\sigma = 1$, Gaussian $\sigma = 10$ and Poissonian plus Gaussian $\sigma = 10$. Therefore, the algorithms were tested in different noisy scenarios (low noise level, Gaussian dominant and Poissonian dominant).

As mentioned in Section 3.1 datasets of 10, 20, 50 and 100 binned images, together with the original 200-frame dataset, were used as input for IDAC to test the trade-off between the (implied) PSF diversity and SNR per frame. We tried to shed some light on the question which of the following strategies is better: deconvolving a long exposure image with a relatively high SNR (AWMLE, ACMLE and MISTRAL algorithms) or tackling the problem by dividing the dataset into more (diverse) frames at the expense of reducing the SNR for each frame (MFBD).

It was found that the original 200-frame set yielded significantly worse results than the smaller sets. Specifically, the 50-frame dataset proved to be the best input to IDAC although the differences between its output and those of the 10 and 20-frame sets were small.

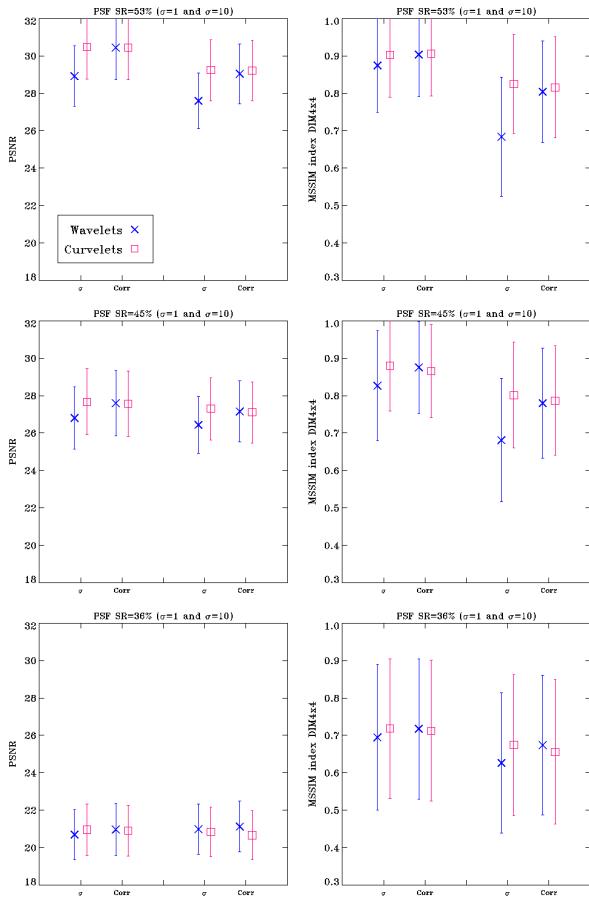


Fig. 8. Left: PSNR index calculated from results obtained in the wavelet and in the curvelet domain with two different types of mask (standard deviation and correlation). Right: 4×4 MSSIM index. Every plot shows results which were obtained from images contaminated with Gaussian noise of level $\sigma = 1$ and $\sigma = 10$. Three levels of PSF mismatch have been tested: top row SR=53%, middle row SR=45%, bottom row SR=36%. Crosses: wavelet-based approach (AWMLE). Squares: curvelet-based approach (ACMLE).

Figure 9 shows the mean value and standard deviation of all the PSNR and 4×4 SSIM windows measured between the reconstructed image and the original one. Under low-noise conditions it is possible to observe that MLE based on either wavelets or curvelets gives best results even for a PSF mismatch of 6-8% in terms of SR. At this stage, we want to stress the apparent contradiction between this statement and the particular case shown in figure 9, middle row, readout noise of $\sigma = 1$, as well as the different behaviour between PSNR and MSSIM indices in this situation. Looking at figure 10, where MSE maps obtained from MISTRAL and ACMLE reconstructions at different scales of representation are shown, one can see that ACMLE presents better values in all the pixels of the object (planet's body, rings, space between the body and the innermost ring) except the pixels at the edges of the main planet body and the rings which show poor numbers. For these pixels, $L_2 - L_1$ object prior implemented in MISTRAL is working impressively well, while ACMLE's values show large departures from the real object. In this situation, SSIM boundedness property, which fixes an upper limit of the metric, offers a better view of the algorithms performance.

When the mismatching of the calibrator is as much as 15-17% the performance of a static-PSF approach deteriorates significantly as shown in figure 9, bottom row, left column. The

quality of reconstructions for the myopic/blind codes is more uniform with respect to the PSF mismatch which is a logical result since such algorithms are designed to deal with large differences between the science PSF and the reference. At the level of mismatch of $\sim 15\%$ there is already a visible qualitative difference between the PSFs (compare the leftmost and rightmost panels of Figure 1) and this affects the performance of static-PSF codes. We again want to point out the contradiction between results shown by PSNR and MSSIM when PSF of SR=36% is used. Photometric values achieved by static-PSF approaches exhibit larger departures but the visual quality is not very degraded when the reconstruction and the real object are compared (in different scales), i.e., local behaviours within each SSIM window do not show significant differences in variance and correlation. In this sense and in this case, the PSNR metric works better. These two aforementioned contradicting results from PSNR and MSSIM demonstrate the need to employ more than one tool to evaluate the performance of deconvolution algorithms because there are many points of view concerning the tasks for which the reconstructions will be used.

When the noise level increases, differences in SR between the PSFs are not as important as the noise level itself. The performances of MFBD and MISTRAL are more affected by noise (Fig. 9, top and middle row). In general, we found that IDAC has poor noise control. The criteria used by IDAC were not able to create a well-balanced trade-off between noise reconstruction and the achieved resolution enhancement, yielding either excessively noisy reconstructions or poor resolution. When comparing AWMLE/ACMLE and MISTRAL in these noisier cases we see that the former still give better results when the SR mismatch is smaller than 10%. These improvements are of the order of 2-4dB in terms of PSNR or 10-20% in terms of MSSIM, depending on the noise level (readout noise of $\sigma = 10$ or both shot and readout noise) and the mismatching between the PSFs (SR=53% or SR=45% cases).

Figures 6, 12 and 13 show reconstructions, PSNR and SSIM maps for the SR=45% calibrator. AWMLE, ACMLE and MFBD exhibit the typical ringing effect associated with strong transitions or edges (very visible at the limits of Saturn's body) whereas MISTRAL's $L_2 - L_1$ edge-preserving prior attenuates such effects considerably. On the other hand, this prior must be responsible for the excessive attenuation of some elongated features present on Saturn's inner ring (Fig. 6, subpanels in row 4), while they have been detected by MFBD, AWMLE and ACMLE (although over-reconstructed in the latter two cases). We stress the fact that such features are not visible in the blurred and corrupted image (Fig. 6, subpanel 4.A). AWMLE and ACMLE exhibit the best results in the planet's body (in terms of the achieved resolution and noise attenuation), and also in the background space between Saturn's body and the innermost ring, which is a measurement of the algorithm's ability to suppress the noise (Figs. 6, 12 and 13, subpanels in row 1). This effect is especially visible in Cassini's division (Figs. 6, 12 and 13, subpanels in row 4). We point out the curious opposite performance of MFBD and static-PSF approaches within Saturn's body: while MFBD achieves its better photometric result in the south pole, AWMLE and ACMLE obtain their better values in the north pole (Fig. 6, subpanels 2.B, 2.D and 2.E).

Amplification of noisy structures degrades the quality of PSNR, and especially, SSIM maps obtained by IDAC, probably to a lesser extent if compared to the typical Richardson-Lucy deconvolution (Fig. 6 subpanel 2.A) but noise reconstruction is absolutely better controlled by the multiresolution support (Columns D and E in Figs. 6, 12 and 13). MISTRAL also

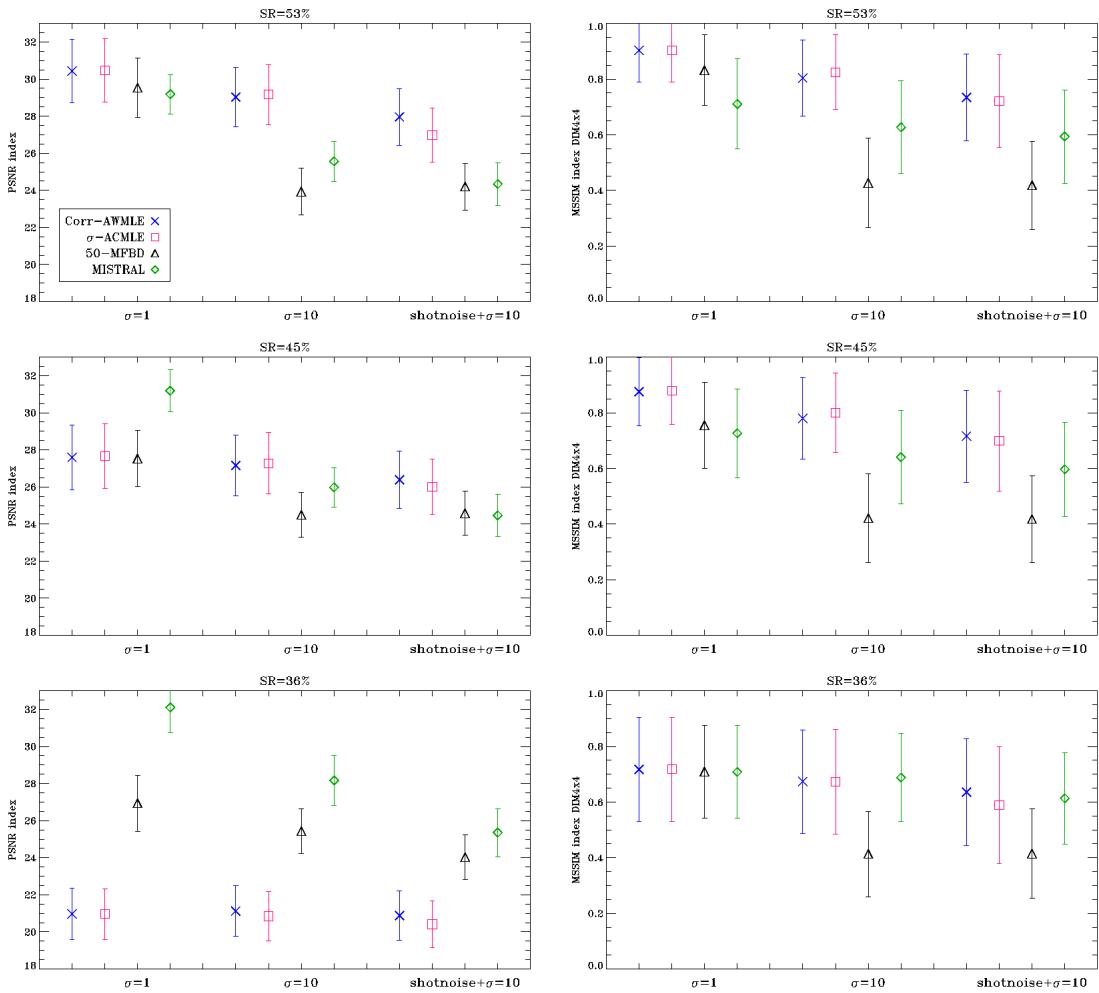


Fig. 9. Each plot shows the mean value and standard deviation of the PSNR and 4×4 MSSIM metrics calculated from results obtained with AWMLE and ACMLE (the former using the correlation-based mask and the latter using the σ -based mask), MFBD applied over a 50-frame dataset, and MISTRAL. Deconvolution was performed with three PSFs in order to represent different levels of miscalibration. Left column: PSNR. Right column: 4×4 MSSIM. Top row: dataset was deconvolved with a matched PSF (SR=53%). Middle row: dataset was deconvolved with mismatched PSF (SR=45%). Bottom row: dataset was deconvolved with highly-mismatched PSF (SR=36%). Crosses: AWMLE with correlation-based mask. Squares: ACMLE with σ -based mask. Triangles: MFBD (50 frames). Diamonds: MISTRAL.

achieves good noise control. However, its performance is not better than that of AWMLE and ACMLE in the body part of Saturn (Fig. 12, subpanels 2,3.C vs. 2,3.D/E and Fig. 13, subpanels 2,3.C vs. 2,3.D/E), whereas MISTRAL's performance is more homogeneous across the rings (Fig. 12, subpanels 6.C vs. 6.D/E and Fig. 13, subpanels 6.C vs. 6.D/E) where AWMLE and ACMLE show poor behaviour, especially when both shot and readout noises are present. On the other hand, photometric accuracy of MISTRAL is worse in the same region (Fig. 12, subpanels 5.C vs. 5.D/E and Fig. 13, subpanels 5.C vs. 5.D/E), thus showing again the complementarity between MSE and SSIM metrics. Finally, ACMLE has been able to reconstruct lines and elongated features within Saturn's body, at high noise levels (Fig. 12 and 13, subpanel 1.E) without the visible dotted texture generated by the rest of algorithms, typical of non-complete elimination of noise.

6. CONCLUSIONS

We have introduced a way of using multiresolution support, applied in the wavelet and curvelet domain, in the post-processing of adaptive-optics images. We have shown how a

correlation-based mask applied over two simultaneously deconvolved images can help control the process of noise amplification. Furthermore, we have introduced the use of the structural similarity index as a measure of performance of different deconvolution algorithms (although it is not our aim in this paper to substitute the necessary subjective visual inspection of the results), and compare its results with those achieved by the more common metric, the mean squared error, showing the need of using more than one metric for comparing the quality of deconvolved images.

We have shown that the use of probabilistic masks can control noise amplification in a Richardson-Lucy scheme but, on the other hand, it does not prevent such amplification indefinitely (Fig. 5). We also have shown that curvelet transform (CT) performs better than wavelet transform (WT) when a σ -based mask is used (Fig. 8).

One of the most important goals of this research was to devise an objective check of the typical assumption (within the AO astronomical community) about the supposed poor performance of static-PSF approaches with respect to the blind/myopic methods. For the dataset we used, the Richardson-Lucy scheme, controlled over the wavelet domain by a correlation mask, and over

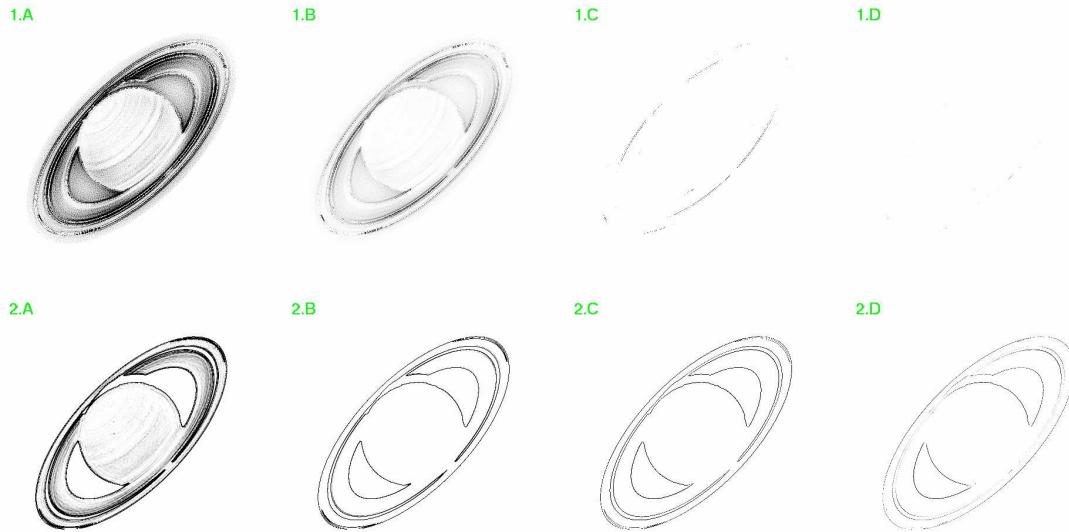


Fig. 10. Row 1: MSE maps of MISTRAL reconstruction. Row 2: MSE maps of ACMLE reconstruction (σ -based mask). Column A: linear scale from 0 to 2000. Column B: linear scale from 2000 to 10000. Column C: linear scale from 10000 to 20000. Column D: linear scale from 20000 to 40000. Reconstructions were obtained from datasets corrupted with Gaussian noise of $\sigma = 1$. Reference PSF at SR=45%. Inverted color scale was used to keep the general choice of representing better qualities with higher brightness. These reconstructions correspond to the particular case shown in figure 9, middle row, left column, $\sigma = 1$, green diamond (MISTRAL) and pink square (ACMLE). Note that averaged PSNR for MISTRAL presents a better value due exclusively to the bad performance of the rest of algorithms in the border of the object with the zero-valued background, whereas ACMLE exhibits a better performance for the rest of pixels.

the curvelet domain by a σ -based mask, provided very competitive performance against the more well-known approaches like MFBD or regularized deconvolution (MISTRAL). Specifically, in the low-noise scenario it yielded 10-15% better results (in terms of MSSIM) than IDAC and MISTRAL for mismatching in the PSF of up to approximately 8% in terms of the Strehl ratio (middle panel, right column in Fig. 9). For higher noise levels the results of AWMLE and ACMLE were still 10-15% better than those of MISTRAL, again up to 8% mismatch, and they were 30-40% better than those of IDAC. In terms of PSNR, the improvement is of 1.5-2dB with respect to MISTRAL and 2-3dB with respect to IDAC. On the other hand, performance of MISTRAL and IDAC was photometrically better when highly mismatched PSF of SR=36% was employed as calibrator.

The observed poor performance of IDAC, which can be visually appreciated in Figures 6-13, can be explained in the following way. Multi-frame approaches rely on diverse PSFs to separate the blurring kernel from the object and to make the blind problem more tractable (less under-determined). Having more images helps when one has truly diverse PSFs. But this is not always the case with AO imaging. In fact, AO will stabilize the PSF and no number of new frames can then supply new information for MFBD. The standard deviation of the Strehl ratio value across the 200 frames used as input for IDAC was only 2%. For all datasets that we had (10 stars) standard deviation across 200 frames rarely exceeded 5%. It seems that single-image codes like MISTRAL or AW(C)MLE have an advantage in the case of stable AO observations.

Several lines of research are still open. It would be possible to study other types of masks, such as those based on the quadratic distance between related zones in consecutive wavelet planes (Starck & Murtagh 1994) or image segmentation (or wavelet plane segmentation) by means of neural networks in order to link and classify significant zones in the image (Núñez & Llacer 1998). Furthermore, the use of other wavelet transforms,

such as the undecimated Mallat trasform with three directions per scale, and many other multitransforms should be studied and compared, e.g., shearlets (Guo & Labate 2007) or waveatoms (Demanet & Ying 2007), to name only two.

Much more important, in our opinion, is studying how to include the multitransform support or probabilistic masks into blind and myopic approaches in order to improve their performance in terms of controlling the noise reconstruction inherent to every image reconstruction algorithm. This will be the topic of our future research.

Acknowledgements. We would like to thank Nicholas Law (University of Toronto) for supplying the PSFs from the Palomar Observatory. We extend our gratitude to Xavier Otazu (Universitat Autònoma de Barcelona), Julian Christou (Gemini Observatory) and Laurent Mugnier (ONERA - The French Aerospace Lab) for useful comments and suggestions about the correct usage of algorithms AWMLE, IDAC and MISTRAL, respectively. We would also like to thank our referee whose suggestions and comments helped us to improve the paper. Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command (USAF) under grant number FA8655-12-1-2115, and by the Spanish Ministry of Science and Technology under grant AyA-2008-01225. The U.S Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

References

- Andrews, H.C. & Hunt B.R. 1977, Digital Image Restoration (Prentice-Hall, Inc., Englewood Cliffs, New York)
- Baena Gallé, R. & Gladysz, S. 2011, PASP, 123, 865
- Baena Gallé, R. & Núñez, J. 2010, IJIST, 20(4), 333
- Bertero, M. & Boccacci, P. 1998, Introduction to inverse problems in imaging (Inst. Phys., London)
- Blanco, L. & Mugnier L.,M. 2011, Optics Express, 19(23), 23227
- Bouman, C. & Sauer, K. 1993, IEEE Trans. Image Proc., 2, 296
- Candès, E., Demanet, L., Donoho, D. & Ying, L. 2005, Multiscale Model. Simul., 5, 861
- Conan, J.-M., Fusco, T., Mugnier, L.M., Kersalé, E. & Michau, V. 1999, Astronomy with adaptive optics: present results and future programs. ESO Conference and Workshop Proceedings, vol. 56, 121 (Bonaccini, D. ed, Garching bei München, Germany)

Demanet, L. & Ying, L. 2007, *Appl. Comp. Harmon. Anal.*, 23(3), 368

Diolaiti, E., Bendinelli, O., Bonaccini, D., Close, L., Currie, D. & Parmeggiani, G. 2000, *A&A Suppl. Ser.*, 147, 335

Esslinger, O. & Edmunds, M. G. 1998, *A&A Suppl. Ser.*, 129, 617

Fadili, M.J. & Starck, J.L. 2006, Sparse representation-based image deconvolution by iterative thresholding (ADA IV, Elsevier, France)

Frieden, B. 1978, *Image enhancement and restoration* (Springer, Berlin)

Gladysz, S., Baena-Gallé, R., Christou, J. C. & Roberts, L. C. 2010, Proc. of the AMOS Tech. Conf., E24

González Audícan M., Otazu, X., Fors, O. & Seco, A. 2005, *Int. Journal of Remote Sensing*, 26, 595

Green, P.J. 1990, *IEEE Trans. Med. Imaging*, 9, 84

Guo, K. & Labate, D. 2007, *SIAM J. Math Anal.*, 39, 298

Hayward, T.L., Brandl, B., Pirger, B., Blacken, C., Gull, G.E., Schoenwald, J. & Houck, J.R. 2001, *PASP*, 113, 105

Hestroffer, D., Marchis, F., Fusco, T. & Berthier, J. 2002, *Research Note A&A*, 394, 339

Holschneider, M., Kronland-Martinet, R., Morlet, J. & Tchamitchian, P. 1989, *Wavelets, Time-Frequency Methods and Phase Space* (Springer-Verlag, Berlin)

Knox, K.T. & Thompson, B.J. 1974, *ApJ*, 193, L45

Jaynes, J. T. 1982, *Proc. IEEE*, 70, 939

Jefferies, S. & Christou, J. C. 1993, *ApJ*, 415, 862

Labeyrie, A. 1970, *A&A*, 6, 85

Lambert, P., Pires, S., Ballot, J. et al. 2006, *A&A*, 454, 1021.

Lane, R.G. 1992, *JOSA A*, 9, 1508

Llacer, J. & Núñez, J. 1990, The restoration of HST images and spectra (STScI, Baltimore)

Llacer, J., Veklerov, E., Coakley, K.J., Hoffman, E.J. & Núñez, J. 1993, *IEEE Trans. Med. Imag.*, 12, no. 2, 215

Lohmann, A.W., Weigelt, G. & Wirnitzer, B. 1983, *Applied Optics*, 22(24), 4028.

Lucy, L.B. 1974, *AJ*, 79, 745

Lucy, L.B. 1994, The restoration of HST images and spectra II (STScI, Boston)

Lühe, O. von der. 1984, *JOSA A*, 1, 510

Matson, C.L. & Haji, A. 2007, Proc. of the AMOS Tech. Conf., E80

Matson, C.L., Borelli, K., Jefferies, S., Beckner, C.C., Hege, E.K. & Lloyd-Hart, M. 2009, *Applied Optics*, 48(1), A75

Molina, R., Núñez, J., Cortijo, F. & Mateos, J. 2001, *IEEE Signal Proc. Magazine*, 18, 11

Mugnier, I.M., Fusco, T. & Conan, J.M. 2004, *JOSA A*, 21, 1841

Núñez, J. & Llacer J. 1993, *PASP*, 105, 1192

Núñez, J. & Llacer J. 1998, *A&AS*, 131, 167

Otazu, X. 2001, *Algunes aplicacions de les wavelets al procés de dades en Astronomia i Teledetecció*, phD thesis, (University of Barcelona, Barcelona), <http://www.tesisenred.net/handle/10803/732>

Otazu, X., González Audícan M., Fors, O. & Núñez, J. 2005, *IEEE Trans. on GeoScience and Remote Sensing*, 43, 2376

Richardson, W.H. 1972, *JOSA*, 62, no. 1, 55

Schulz, T.J. 1993, *JOSA A*, 10(5), 1064

Shensa, M.J. 1992, *IEEE Trans. on Signal Proc.*, 40(10), 2464-2482

Starck, J. L. & Murtagh, F. 1994, *A&A*, 288, 342-348

Starck, J. L., Murtagh, F., Querre, P. & Bonnarel, F. 2001, *A&A* 368, 730

Starck, J. L., Donoho, D.L. & Candès, E. 2002a, *A&A*, 398, 785

Starck, J. L., Pantin, E. & Murtagh F. 2002b, *PASP*, 114, 1051

Starck, J. L., Aghanim, N. & Forni, O. 2004, *A&A* 416, 9

Starck, J. L., Murtagh, F. & Fadili, M.J. 2010, *Sparse Image and Signal Processing: wavelets, curvelets, morphological diversity* (Cambridge University Press, Cambridge)

Storrs, A. & Eney, B. 2010. *Bulletin Amer. Astr. Soc.*, 42, 456

Thiébaut, E. & Conan, J.M. 1995, *JOSA A*, 12(3), 485

Tikhonov, A.N., Goncharski, A.V., Stepanov, V.V. & Kochiko, I.V. 1987, *Sov. Phys. Doklady*, 32, 456

Trauger, J.T., Ballester, G.E., Burrows, C.J. et al. 1994, *ApJ*, 435, L3

Troy, M., Dekany, R.G., Brack, G.L. et al. 2000, *Proc SPIE*, 4007, 31

Wang, Z., Bovik., A.C., Sheikh, H.R. & Simoncelli, E.P. 2004, *IEEE Trans. Im. Proc.*, 13(4), 600

Wang, Z. & Bovik., A.C. 2006, *Modern image quality assessment* (Morgan & Claypool, ebook)

Worden, S.P., Stein, M.K., Schmidt, G.D. & Angel, J.R.P. 1977, *Icarus*, 32, 450.

Zellner, N.E.B., Gibbard, S. & de Pater, I. 2005, *Icarus*, 61, 355

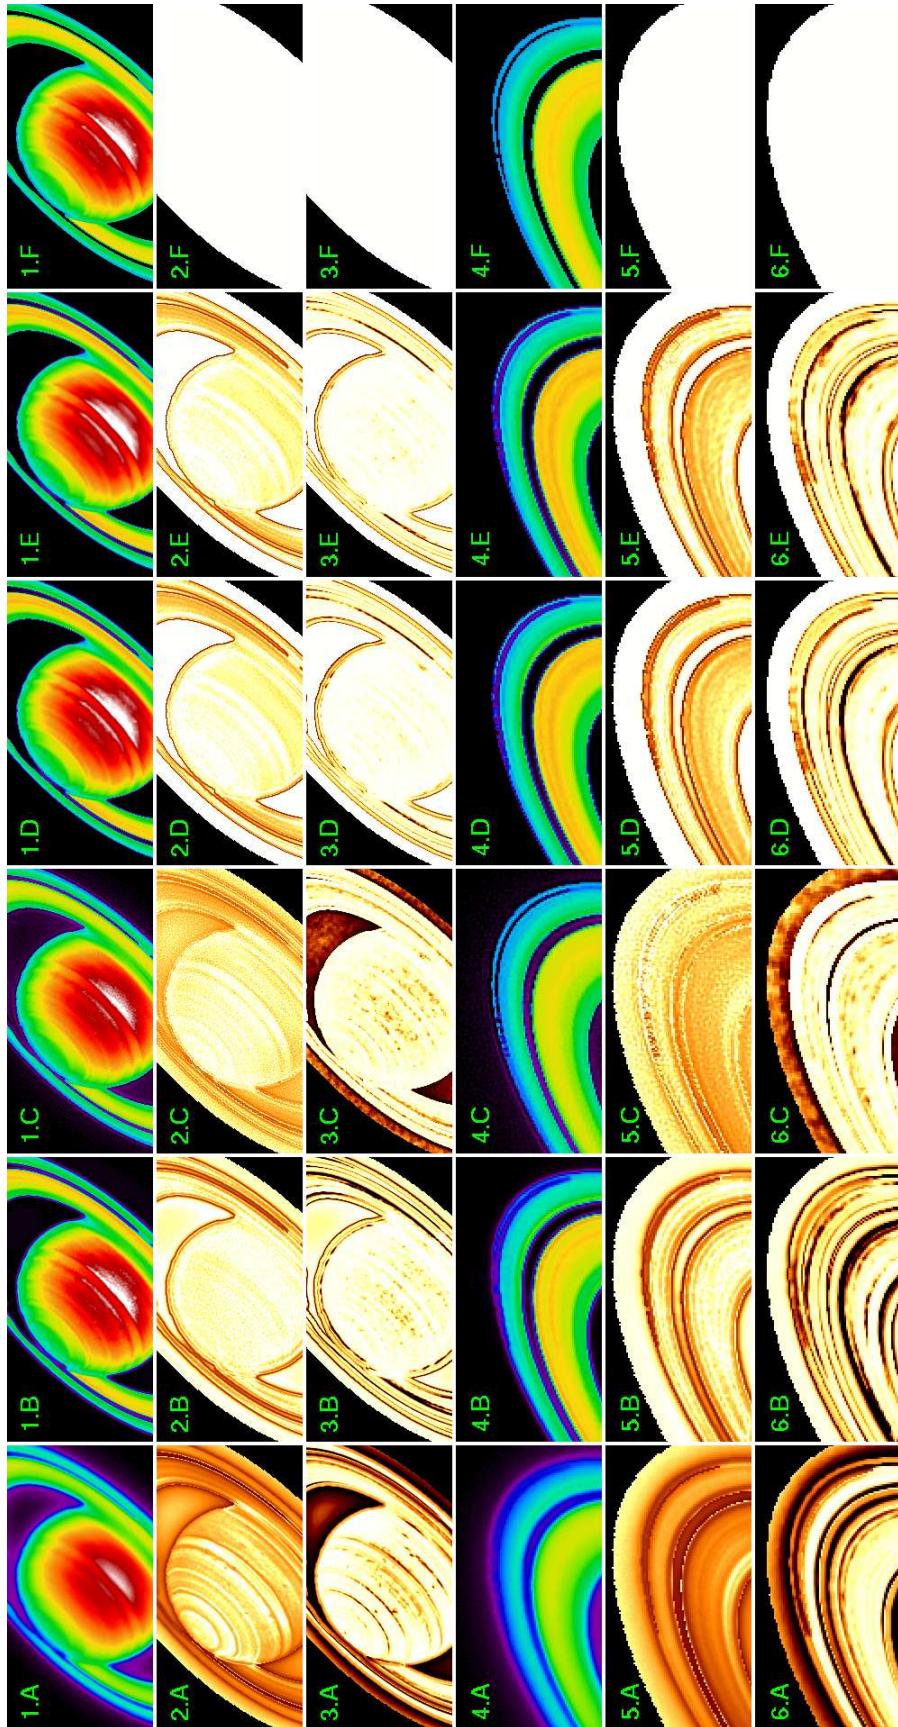


Fig. 11. Row 1: Saturn's body; 1.A: Degraded and corrupted image with Gaussian noise ($\sigma = 1$), 1.B: MFBD reconstruction (50 frames), 1.C: MISTRAL reconstruction, 1.D: AWMLE reconstruction (correlation-based mask), 1.E: ACMLE reconstruction (σ -based mask), 1.F: Original image. Row 2: corresponding PSNR maps. Row 3: corresponding 4×4 SSIM maps. Row 4: Saturn's rings detail. Row 5: corresponding PSNR maps. Row 6: corresponding 4×4 SSIM maps. Reconstructions were performed with reference PSF at SR=45%. Rows 1 and 4 are represented on a linear scale from 0 to 975, matching the dynamic range of the original image, i.e., subpanels 1.F/4.F. Rows 2 and 5 are represented on a linear scale from 0 to 50. Rows 3 and 6 are represented on a linear scale from 0 to 1. In rows 2, 3, 5 and 6 higher brightness indicates better quality.

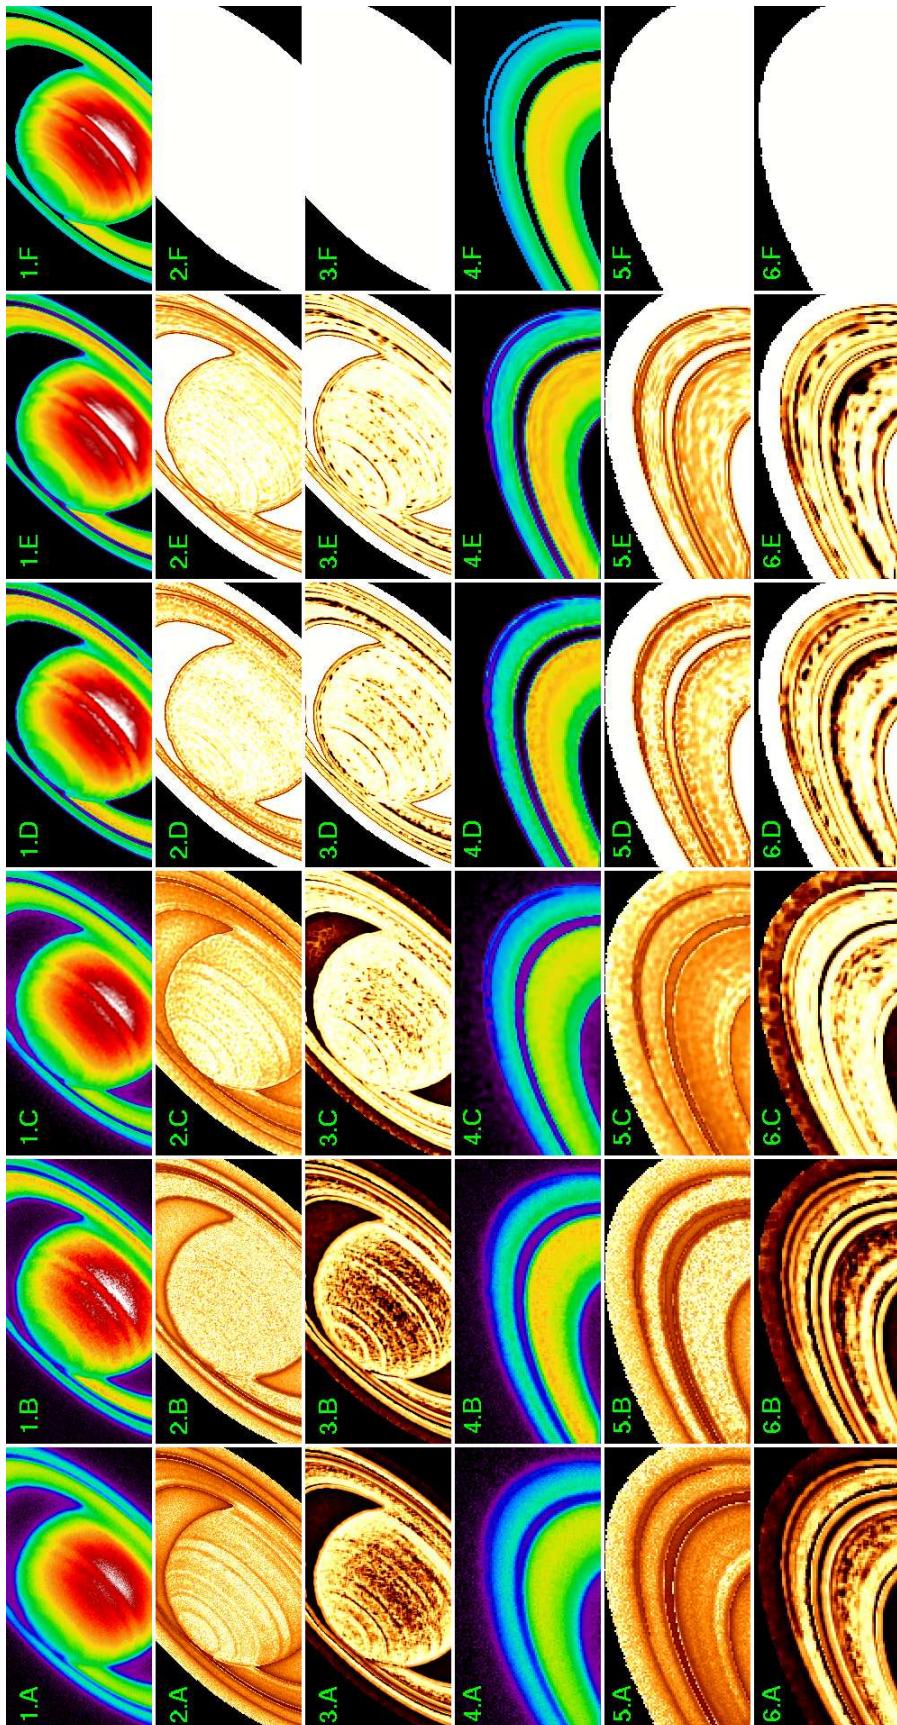


Fig. 12. Row 1: Saturn's body; 1.A: Degraded and corrupted image with Gaussian noise ($\sigma = 10$), 1.B: MFBD reconstruction (50 frames), 1.C: MISTRAL reconstruction, 1.D: AWMLE reconstruction (correlation-based mask), 1.E: ACMLE reconstruction (σ -based mask), 1.F: Original image. Row 2: corresponding PSNR maps. Row 3: corresponding 4×4 SSIM maps. Row 4: Saturn's rings detail. Row 5: corresponding PSNR maps. Row 6: corresponding 4×4 SSIM maps. Reconstructions were performed with reference PSF at SR=45%. Rows 1 and 4 are represented on a linear scale from 0 to 975, matching the dynamic range of the original image, i.e., subpanels 1.F/4.F. Rows 2 and 5 are represented on a linear scale from 0 to 50. Rows 3 and 6 are represented on a linear scale from 0 to 1. In rows 2, 3, 5 and 6 higher brightness indicates better quality.

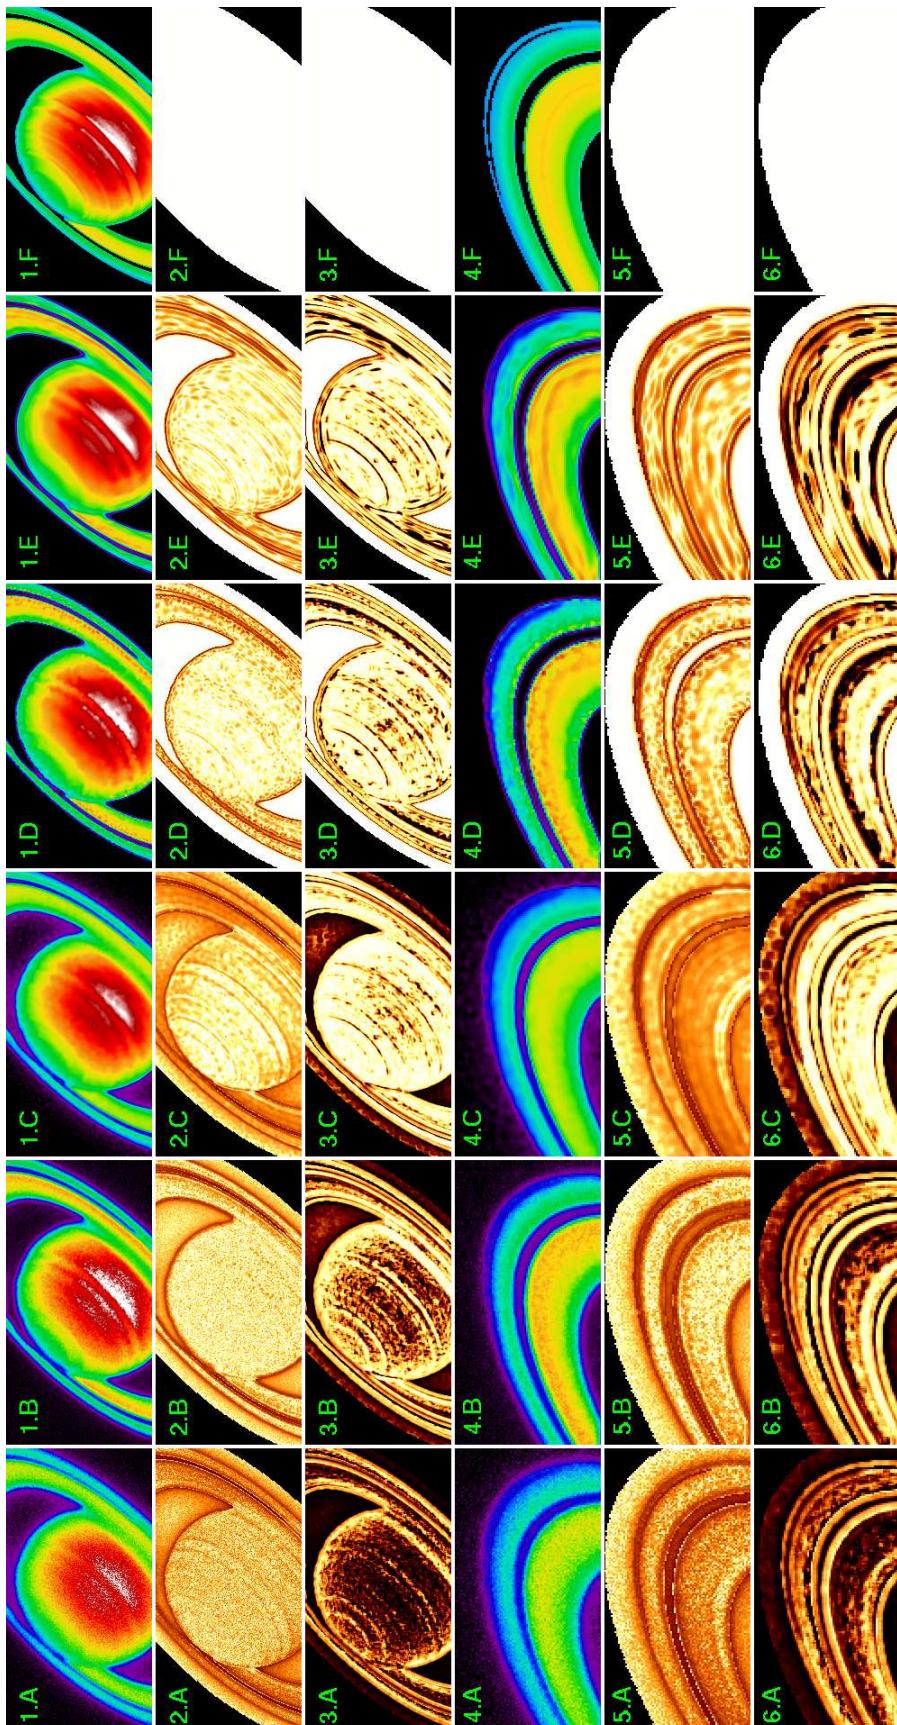


Fig. 13. Row 1: Saturn's body; 1.A: Degraded and corrupted image with both shot and Gaussian noise ($\sigma = 10$), 1.B: MFBD reconstruction (50 frames), 1.C: MISTRAL reconstruction, 1.D: AWMLE reconstruction (correlation-based mask), 1.E: ACMLE reconstruction (σ -based mask), 1.F: Original image. Row 2: corresponding PSNR maps. Row 3: corresponding 4x4 SSIM maps. Row 4: Saturn's rings detail. Row 5: corresponding PSNR maps. Row 6: corresponding 4x4 SSIM maps. Reconstructions were performed with reference PSF at SR=45%. Rows 1 and 4 are represented on a linear scale from 0 to 975, matching the dynamic range of the original image, i.e., subpanels 1.F/4.F. Rows 2 and 5 are represented on a linear scale from 0 to 50. Rows 3 and 6 are represented on a linear scale from 0 to 1. In rows 2, 3, 5 and 6 higher brightness indicates better quality.